

Translating Across Cultures: LLMs for Intralingual Cultural Adaptation

Pushdeep Singh, Mayur Patidar, Lovekesh Vig
TCS Research

{pushdeep.singh, patidar.mayur, lovekesh.vig}@tcs.com

Abstract

LLMs are increasingly being deployed for multilingual applications and have demonstrated impressive translation capabilities between several low and high-resource languages. An aspect of translation that often gets overlooked is that of cultural adaptation, or modifying source culture references to suit the target culture. While specialized translation models still outperform LLMs on the machine translation task when viewed from the lens of correctness, they are not sensitive to cultural differences often requiring manual correction. LLMs on the other hand have a rich reservoir of cultural knowledge embedded within its parameters that can be potentially exploited for such applications. In this paper, we define the task of cultural adaptation and create an evaluation framework to evaluate the performance of modern LLMs for cultural adaptation and analyze their cross-cultural knowledge while connecting related concepts across different cultures. We also analyze possible issues with automatic adaptation. We hope that this task will offer more insight into the cultural understanding of LLMs and their creativity in cross-cultural scenarios.

NOTE: This paper contains examples that may be offensive.

1 Introduction

Recent progress in NLP is largely driven via LLMs, which have shown great promise in a variety of tasks including text generation, language understanding, question answering, code generation, and even machine translation. Though LLMs have not achieved state-of-the-art performance for machine translation (Zhu et al., 2023), their instruction-following ability makes them suitable for tasks involving more creativity and customization during generation. Many translation applications require literal translations for which specialized transformer-based models trained on parallel data

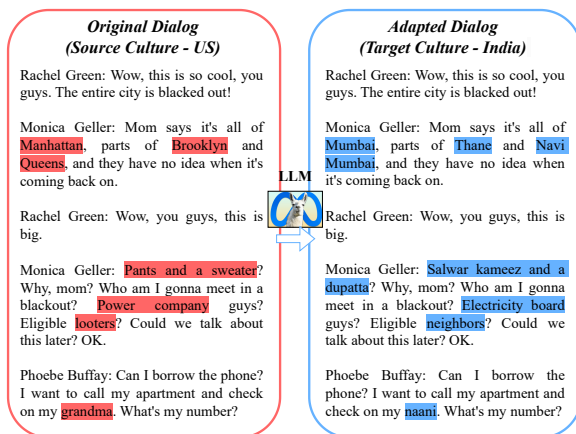


Figure 1: Cultural Adaptation using LLM

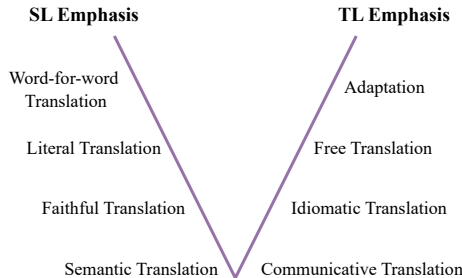


Figure 2: Newmark (1988)'s V diagram of translation methods. SL: Source Language, TL: Target Language

are ideal. However, there are other facets of translation (see Figure 2), such as adaptation, also called the 'freest' form of translation (Newmark, 1988) wherein the original text is rewritten to make it more appropriate for the target audience belonging to a specific age group or culture (See Figure 1). Applications of adaptation (Appendix E) include adapted transcriptions for plays, poetry, and movie subtitles where the plot, characters and central theme are usually kept intact but the text is rewritten to ensure the output is sensitive to the target culture. Adaptation can either be done within the same (intralingual adaptation) or in different languages (interlingual adaptation).

Polizzotti (2018) in his book “*Sympathy for the Traitor: The Translation Manifesto*” describes how in 17th century France, a sexist term *belles infidèles* (the beautiful, unfaithful ones) was used to describe the prevalent approach to French translations at the time, which involved “updating” ancient Greek and Latin texts by removing vulgar language or sexual content and replacing outdated references with modern equivalents to make the texts more easily understandable and socially acceptable. These translations were considered “beautiful” because they were smooth to read and met contemporary expectations, but they were not faithful to the original texts in a strict sense. The debate between “literalism” and “adaptation” persists, with proponents of each arguing their merits. Yet, adaptations of existing texts continue to serve diverse purposes including cross-cultural communication.

In this study, we steer clear of this debate and explore this task purely from an NLP perspective particularly investigating the power of large language models. We define a specific version of the task along with clear goals and an evaluation framework for assessing the effectiveness of these adaptations considering factors such as localisation, preservation, naturalness, and appropriateness. The motivation behind this work stems from the need to transcend the constraints of literal translation and explore freer forms of translation such as adaptation. Due to the rising creativity, multilinguality, cross-cultural knowledge and instruction-following ability of modern language models, they have the potential to generate culturally resonant adaptations of the source text.

We limit our study to cultural adaptation with English as the source and target language i.e. Intra-lingual adaptation. As Hershcovich et al. (2022) argues, although language and culture are interconnected, they are not synonymous. For example, English, being the *lingua franca* for many parts of the world, can carry views and concepts from different parts of the world. By sticking to English, we can specifically evaluate how well cultural aspects are adjusted in adaptation without the added complexity of translating between languages. As LLMs become more multilingual (in generation and understanding), their ability can better be evaluated for interlingual adaptation and related aspects of this study can be applied there. We can also view Interlingual Cultural Adaptation as a combination of Intra-lingual Cultural Adaptation and Machine Translation.

We explore the following research questions and contribute along these: RQ 1) How do we define what constitutes adaptation in terms of modifications to the source text i.e. what is changed during adaptation and for what purpose? RQ 2) Based on the goals of adaptation, what are the optimal criteria/aspects for evaluation? RQ 3) Given the evaluation, how proficient are modern language models at adaptation? What strategies do they employ, and to what extent do they adapt based on provided instructions? RQ 4) What insights does this offer into their parametric cross-cultural knowledge?

2 Related Work

Yao et al. (2023) discusses the aspect of using cultural knowledge to support LLM-based translation. They focus on literal translation and create a culture-specific parallel corpus, to evaluate the cultural awareness of MT systems. They explored different prompting strategies using external and internal knowledge for LLM-based machine translation and created an automatic evaluation metric, to measure the translation quality of cultural concepts.

Recent works on evaluating cultural awareness in LLMs have centred primarily around measuring cultural value alignment (Durmus et al., 2023; Cao et al., 2023; Masoud et al., 2023; Ramezani and Xu, 2023). While this is important, it does not necessarily indicate that LLMs are aware of culture-specific items or concepts from different cultures. More research is needed to assess whether LLMs truly understand these culture-specific items and concepts and can use them coherently in text. Our research aims to address this question.

Jiang and Joshi (2023) created a ranking-based statistical QA task that compared cultural concept popularity across countries. Wang et al. (2024) examined the cultural dominance of concrete (e.g., holidays and songs) and abstract (e.g., values and opinions) cultural objects in LLM responses.

Peskov et al. (2021) introduced the idea of automatic cultural adaptation by adapting named entities across cultures and languages, however, it focused on simpler entities in standalone sentences. Cao et al. (2024) constructed resources for cultural adaptation of recipes and also evaluated their method against LLM-based adaptation. More recently, Zhang et al. (2024) created Chinese-English menu corpora and defined an evaluation for the task

of adapting restaurant menus.

3 Task Definition

For the task of adaptation, we use a corpus of dialogs from a TV show and adapt it to the target culture. We choose adaptation of dialogs instead of standalone sentences as done by Peskov et al. (2021) since they provide richer context and are more representative of a true use case of adaptation. The original corpus of dialogues is denoted as $D_o = \{d_1^o, d_2^o, \dots, d_n^o\}$. We obtain an adapted version of these dialogues denoted as $D_a = \{f(d_1^o, c), f(d_2^o, c), \dots, f(d_n^o, c)\}$, where f represents the language model that adapts the original dialogues to the target culture and c is the specific cultural context or prompt representing the cultural context for adaptation. Each dialogue d consists of a number of utterances i.e. $d = \{u_1, u_2, \dots, u_m\}$. Each utterance $u_i = \text{speaker}(u_i) : \text{text}(u_i)$, where $\text{speaker}(u_i)$ is the speaker or participant name and $\text{text}(u_i)$ is the textual content for utterance u_i . Our task is to evaluate how well dialogues in the adapted set D_a are culturally aligned to the target culture while maintaining the intent and essence of the original dialogue. Section 4 provides details on the exact aspects along which we assess these adaptations.

4 Annotating Cultural References

Corpus Description: We choose the ‘Friends Dialogs’ corpus for this study. We filter the data to choose dialogues with utterances between 1 and 15. The corpus includes 1110 conversations (or dialogs) containing 11812 utterances by 363 speakers. The reason for choosing such a corpus is that ‘Friends TV Show’ is deeply rooted in American culture offering a distinct contrast that highlights the need for adaptation when targeting a new cultural context, specifically *India*¹ in our study. Here, adaptation ensures that the message is not only understood but also embraced and valued in a different cultural environment.

Culture is a multi-faceted concept. Many scholars have tried to define culture. One such theory which is very relevant here and is also mentioned in cultural translation studies is Hall’s

¹We choose country as a proxy for culture (Adilazuarda et al., 2024). While a country such as India has many subcultures, still, many aspects and items are still universal and are relatable to a national audience. Those remain the key focus of our study and our annotation task.

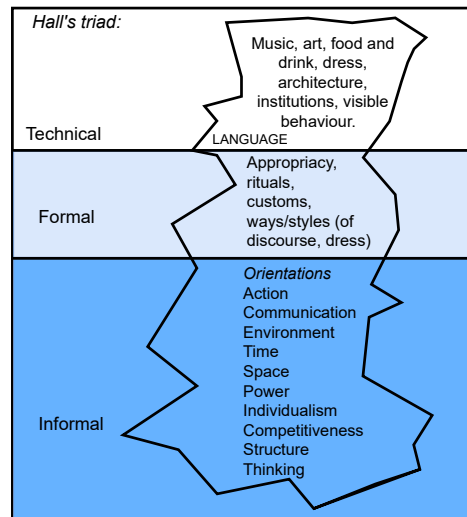


Figure 3: Hall’s Iceberg Theory and Triads

Iceberg Model of Culture (or the Triad of Culture) which divides aspects of culture into three levels: visible (above the waterline), semi-visible and invisible (see Figure 3) which are referred as the technical, formal and informal level of culture, respectively. As Katan (2014) describes, these levels also relate to how we grasp culture: technical culture can be taught by an expert, formal culture through trial-and-error while informal culture is learned unconsciously.

At the tip of the iceberg i.e. the technical level, the goal of translation is to transfer the terms and concepts of the source text to the target text with minimal loss. The terms and concepts are usually referred to as ‘culture-bound’ terms, or ‘culturemes’. Hall’s second level, i.e., the Formal level of culture focuses on rituals, customs, and accepted or appropriate ways of doing things. This level follows the ‘Skopos Theory’ (Vermeer, 1989) i.e. translation should be oriented towards achieving the desired function in the target culture, rather than being faithful to the source text. Hall’s third level, i.e., the Informal level cannot be taught or learned but is acquired ‘out-of-awareness’ or unconsciously. This is what makes a translation more artistic rather than mechanistic.

RQ 1: How do we define what constitutes adaptation in terms of modifications to the source text i.e. what is changed during adaptation and for what purpose? In this study, we mainly focus on the first two levels of culture. In order to evaluate whether an adaptation navigates different levels of culture, we need to annotate culture-related references in the source text and look at

how they are being adapted in the corresponding adaptation. We call these items adaptable items or **Culture-Specific Items(CSI)** used by [Newmark \(1988\)](#).

Items which can undergo adaptation include references to concepts and realities which are foreign to the target culture, socially sensitive and taboo topics, colloquialisms, slang, idioms, figures of speech, humour, or content which can be considered offensive in the target culture. We manually annotate these items in our corpus of dialogues. We also categorise these items into the following categories : 1) Ecology ((flora, fauna, winds etc.) 2) Material Culture (artefacts, food, clothes, houses, towns, transport etc.) 3) Social Culture (work and leisure) 4) Institutions, Organizations and ideas (political, social, religious, social, artistic, administrative, ideas etc.) 5) Gestures and Habits (name of regular behaviour and movements), as proposed by [Newmark \(1988\)](#). Additionally, we introduce four more categories which reflect the need for adaptation: 6) Slang or Figure of Speech, 7) Offensive Content, 8) Socially Sensitive or Taboo Topics and 9) Humour (Since 'Friends' is a sitcom). We use the descriptions from [Newmark \(1988\)](#) plus descriptions of the other four categories as our annotation guidelines.

While [Yao et al. \(2023\)](#) demonstrates an automated approach to annotating culture-bound items, however, for our use case, we observed that it only identifies a fraction of items which can undergo adaptation. Also, CSI are culture-specific not due to their origin but also due to their foreignness to the target culture. For example, sausage is common in both the USA and the UK but still foreign in Indian culture, so manually annotating these items based on the foreignness to the target culture is desired. This is especially important due to the "McDonaldization of Society" ([Ritzer, 1996](#)) where cultural boundaries are becoming blurred and the notion of foreignness is constantly evolving due to migration and cultural exchange. Therefore, we also annotate the degree of foreignness to the target audience to provide a more accurate depiction and expectation since items that are more foreign to the target culture should be more likely to undergo adaptation. We define 3 foreignness levels: 1,2 and 3 for our annotation. Foreignness level 1 consists of items/behaviours which have traceable foreign origin/usage however they are common (in terms of familiarity, integration and perception) in the target culture. For example, pizza, chocolate,

cricket and coffee are fairly common in India. We omit items in foreignness level 1 from our analysis. For items with foreignness level 2, they may be recognized in the target culture, but their usage or significance is somewhat foreign or less familiar. For example, sushi, tacos, k-pop and beer are not very common and not fully assimilated or mainstream in India. Items in foreignness level 3 are largely unfamiliar or perceived as distinctly foreign within the target culture. Some examples include *kimono*, rodeo, thanksgiving etc. which are largely unknown to the Indian audience.

Human annotation: We hired three human annotators from India for our study (both annotation and human evaluation (Section 7.2)), who were able to understand different aspects and sensitivities of Indian culture expressed through English. The annotators were aware of the source ("USA" as a proxy) culture and good at identifying what aspects of it are foreign to India and to what extent. They were instructed to annotate² for these cultural items in the corpus, their categories and foreignness level.

Recent studies ([Schaeckermann et al., 2018](#); [Drapeau et al., 2016](#)) have indicated that deliberation can enhance the quality of answers and even a small number of debates can outperform the wisdom of large crowds ([Navajas et al., 2017](#)). Therefore, in this study, annotations were carried out using deliberation through verbal discussion until a consensus was reached.

Some examples of CSI for different categories or foreignness levels are given in Table 1. The number of occurrences of these CSI for each category/foreignness is shown in Figure 4. The corpus is publicly accessible³.

5 Evaluation of Cultural Adaptation

In order to design an evaluation framework for adaptation, we need to understand the motivation and goals behind it. In the following section, we mention some goals of adaptation which are aspects along which we assess the quality of cultural

²Although we acknowledge subjectivity in terms of annotation on aspects like foreignness, offensive content, taboo topics, etc., the instructions for annotations were made as specific and unambiguous as possible. Annotators were asked to consider a wider target audience to avoid any personal bias when annotating these cultural items in the corpus, their categories, and their level of foreignness.

³<https://github.com/iampushpdeep/CulturalAdaptEval>

CSI Category	CSI Examples
Ecology	sage branches, Vail, Alps, Grand Canyon, San Diego Zoo, Capuchin etc.
Material Culture	meatball sub, MonkeyShine Beer poster, hamburger, Soap Opera Digest etc.
Social Culture	Another World, Thanksgiving, Days of Our Lives, bridesmaids, Halloween etc.
Institutions, Organisations and Ideas	Alan Alda, Mattress King, Wendy's, FICA, Fortunata Fashions etc.
Gestures and Habits	"You licked and you put", "honk honk", "Cha-ching", "step-ity step and jazz hands" etc.
Offensive Content	"go to hell", dumb ass, bitch, "climb out of my butt", "third nipple" etc.
Socially Sensitive and Taboo topics	porn, naked, lust, have sex, undressing etc.
Humour	knock-knock jokes, "get him something like a wrecking ball, or a vile of smallpox" etc.
Foreignness Level	CSI Examples
2	Christmas, Superman, cheesecake, wok, "spill coffee grounds", Porsche etc.
3	graham cracker, Archie and Jughead Double Digest, barca lounge, Swing Kings etc.

Table 1: Examples of CSI along different categories/foreignness levels found in 'Friends' Corpus.




Figure 4: Number of Occurrences of CSI by a) Category, b) Foreignness level. A total of 3192 occurrences were found.

adaptation.

5.1 Aspects of Evaluation

RQ 2: Based on the goals of adaptation, what are the optimal criteria/aspects for evaluation?

Adaptation can be used for a variety of applications⁴ and the goals will vary for each application be it marketing, children's literature, or creative content translation. However, the main goal of adaptation is to serve the target audience, even if it means being unfaithful to the original text. This means that one of the goals is to achieve a shift in cultural levels to make the text more familiar and appropriate to the target culture by adapting more items in the source text. The greater the **Extent of Cultural Adaptation or Localisation**, the higher the chances it will be accepted by the target

⁴In this study, we are exploring the creative side of adaptation however for more serious applications like adapting legal or medical content, factuality is the most important aspect of evaluation which these language models may not guarantee.

audience. Another obvious goal of adaptation is **Cultural Appropriateness and Sensitivity** i.e. respecting the sensitivities of the target culture without being *offensive* and avoiding propagation of harmful *stereotypes*. Sometimes, items adapted to the target culture may not fit well or might appear forced or unnatural. Thus, another goal is **Naturalness** i.e. that adaptation must appear natural and coherent. Changes done to the source text should not disrupt the flow of the text. As mentioned earlier, adaptation used for plays, poetry, drama etc keeps the characters and the central theme intact and only modifies cultural elements. This means that another goal of adaptation is **Content Preservation**. We want adaptation to preserve the original meaning and intent of the dialogue and it should not distort the main message. Since we are dealing with intralingual cultural adaptation, it is very similar to text style transfer (Mir et al., 2019) which also uses metrics like style transfer intensity (in our case, extent of cultural adaptation), content preservation and naturalness. In order to evaluate cultural adaptations along these aspects, we perform two types of analysis: 1) *edit level analysis*, 2) *dialogue level analysis*.

5.2 Edit Level Analysis

% CSI edited: We define a proxy metric to measure the extent of cultural adaptation i.e. % Of CSI edited. We use the annotations in our source corpus and using fuzzy string matching⁵, we calculate what percentage of cultural elements that we have annotated in source text also appear in translation. If they do that means they aren't adapted/edited. Using this we can calculate % CSI edited as :

$$\% \text{ CSI edited} = 100 - \% \text{ of CSI found in adaptations}$$

This metric is not very informative of the

⁵<https://github.com/seatgeek/thefuzz>

quality of edits performed on items or whether that edit was correct or appropriate, however, it does quantify the extent of change or adaptation. We also report %CSI edited for each category and foreignness level.

Aspect Evaluation: For aspect-based evaluation at the edit level, we rate each individual edit on 3 aspects: localization, correctness in context and offensiveness. During adaptation, many edits are also performed on items which are culturally neutral. This is usually done to make the text more localized by creating new cultural items. Therefore we need to identify all edits whether they are on CSI or non-CSI. While there are libraries which can help to identify edits, we observed that LLMs are more suitable for such a task given their language understanding ability. We use **Mixtral**⁶(Jiang et al., 2024) for automatic evaluation in our experiments including identifying edits in each utterance for all the dialogues. Then we ask the LLM to rate each edit on 3 aspects : 1) *Correctness* (0 or 1), 2) *Localization* (0 or 1 or 2), 3) *Offensiveness*(0 or 1) The prompts used for obtaining edits and rating them on these aspects are given in Appendix B. These aspects somewhat relate to the aspects described in Section 5.1. Correctness relates to Naturalness, Localization to Extent of Cultural Adaptation and Offensiveness to Cultural Appropriateness. However, it’s important to note that aspect evaluation at the edit level may not account for the entire context of the dialog but for the edit and the context in which it is used. Once we obtain these ratings, we can analyze and compare adaptations from different LLMs in terms of 1) percentage of correct edits, 2) Average edit localization score and 3) percentage of offensive edits.

Translation Strategies: For each edit corresponding to source culture CSI, we determine the strategy used for adaptation. According to Davies (2003), the following strategies can be used while translating CSI: 1) Preservation, 2) Addition, 3) Omission, 4) Localisation, 5) Globalisation, 6) Transformation and 7) Creation. The prompt for determining the translation strategy for a given CSI edit is given in Appendix B, which also contains a description of these strategies. In the

⁶<https://huggingface.co/mistralai/Mixtral-8x7B-Instruct-v0.1>

context of intralingual adaptation, ‘preservation’ corresponds to no edit. ‘Creation’ corresponds to edits where non-CSI are edited to CSI. We perform the analysis for CSI edits and classify them based on the strategy used.

5.3 Dialog Level Analysis

For dialog level analysis, we directly ask the LLM to rate the adapted dialog given the original dialog on a scale of 1 to 5, along five aspects : 1) Localization, 2) Naturalness, 3) Offensiveness, 4) Content Preservation and 5) Stereotypical behavior all of which fall under the aspects/goals of adaptation described in Section 5.1. The prompt used to score the adapted dialogs is given in Appendix B. We report average aspect scores over all dialogues.

6 Prompting for Cultural Adaptation

In this study, based on our goals, we use a simple prompt which includes our goals and exemplars to guide the LLM for expected adaptations of the dialogs. The adaptation prompt is given in Table 2.

You have to adapt the given dialogue to align with Indian culture and audience while keeping the response in English. Adapt culture-specific references/items (do not change character names) which are foreign to **Indian culture** to align with Indian cultural context, norms, and sensitivities, while maintaining the correctness, coherence and keeping original intent intact. Also adapt very foreign humour, slang or figure of speech unfamiliar to Indian English audiences, offensive and socially sensitive or taboo content while making sure that the intensity of emotions like humour don’t get affected. Ensure that code-mixing is avoided, and output remains in English. Every utterance in the original dialogue should have a corresponding utterance in the adapted version, don’t add or delete utterances or don’t change speakers.

{ 2 shot example (In Appendix C) }

What is the adapted version for the following dialogue :
{ Dialog }

Table 2: Prompt for getting cultural adaptation

We experimented with a few prompts on a small scale to finally select the prompt for this study. We opine that the correct prompt can unlock certain dimensions and improve creativity, however, a detailed study involving large-scale experiments with different prompts would involve working at a deeper level of culture (especially the informal level) and evaluating related aspects as described in Section 4, which is beyond the scope of the present study.

Models: We explore adaptations obtained from




Figure 5: Percentage of CSI edited in a) total, b) along different categories and c) foreignness level.

3 LLMs : Llama-2 70B⁷(Touvron et al., 2023), Llama-3 8B⁸ and Llama-3 70B⁹. These models are state-of-the-art open source LLMs and are cheaply available for inference.

Examples of adaptation from these models for a given dialog and several utterances are given in Appendix D (Table 12 and Table 14).

7 Results and Analysis

RQ 3: Given the evaluation, how proficient are modern language models at adaptation? What strategies do they employ, and to what extent do they adapt based on provided instructions?

7.1 Edit level Analysis

% CSI Edited As shown in Figure 5, %CSI Edited is lowest for Llama-2 70B (45%). This suggests a lower extent of adaptation. Llama-3 8B (82.8%) and Llama-3 70B (79.7%) seem to perform equally well in editing CSI. Items from the ‘Ecology’ category have the highest percentage of items edited due to items that are easier to edit. A higher fraction of items with foreignness level 3(very foreign) were edited compared to items with foreignness level 2 indicating that adaptation is prioritizing changing of more foreign items preferably into localized, more relatable items/expressions.

Aspect level evaluation : The results for aspect level evaluation are given in Table 3. Using Mixtral,

⁷<https://huggingface.co/meta-llama/Llama-2-70b-chat-hf>

⁸<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

⁹<https://huggingface.co/meta-llama/Meta-Llama-3-70B-Instruct>

Aspect	Llama-2 70B	Llama-3 8B	Llama-3 70B
# Edits	3256	12177	5747
Correctness(%)	99.05	99.79	99.87
Localisation(Average)	1.52	1.55	1.74
Localisation(%(0,1,2))	0.5, 46.7, 52.8	0.1, 54.5, 45.4	0, 27.6, 72.4
Offensiveness(%)	0.43	0.30	0.00

Table 3: Edit level scores




Figure 6: Translation strategies used for adapting CSI by percentage for different models

we extracted significant edits (edits causing significant token change) from each adapted dialogue at an utterance level. Llama-3 8B, surprisingly, uses a large number of edits to get the adaptation, compared to other models. All models used in our evaluation have a high percentage of correct edits, the highest for Llama-3 70B, followed by Llama-3 8B and Llama-2 70B. Also, Llama-3 70B exhibits the highest average localization score per edit, primarily due to a larger proportion of edits being highly localized (score 2). Furthermore, Llama-3 70B displays no instances of offensive behaviour in our evaluation. Based on the edit-level analysis, Llama-2 70B performs slightly worse than other models. Nevertheless, LLMs are prone to perform incorrect edits, examples of which are given in Appendix D (Table 13). As shown, LLMs often struggle with instances that involve understanding and reasoning about cultural objects.

Translation Strategies used : We also obtain the type of strategy used for adapting CSI. Since the adaptation is intralingual, ‘preservation’ is already in use whenever CSI is not adapted, as captured by %CSI Edited. Figure 6 shows strategies used by percentage. We observe similar behaviour across all models: most percentage of CSI edits using localization, followed by transformation and then very closely, globalization. Addition and omission-related edits are very rare during adaptation. Some examples of CSI edits and the corresponding translation strategy used are given in Table 4.

Edit	Strategy Used
sexually → romantically	globalisation
Jimmies → tamarind chutney	transformation
Poulet → Dhoni	transformation
FICA → Income Tax	localisation
predicament room → waiting lounge	globalisation
"Son of a bitch" is back → he is back	omission
Wendy's → Haldiram's	localisation
gumball ring → gumball ring. It's not even a real diamond!	addition

Table 4: Examples of extracted CSI edits and the translation strategy used

Aspect	Llama-2 70B	Llama-3 8B	Llama-3 70B
Localisation	3.53	4.36	4.44
Naturalness	4.32	3.97	4.05
Content Preservation	4.56	4.03	4.27
Offensiveness	1.01	1.01	1.00
Stereotypical	1.18	1.62	1.37

Table 5: Dialog level scores

7.2 Dialog level Analysis

Aspect level evaluation : Aspect level scores are shown in Table 5. We report average aspect scores over all the dialogs. In terms of localization, Llama-3 70B clearly outperforms other models. Llama-2 70B performs the worst in terms of localization, which was also indicated by a lower %CSI edited number as observed in Section 7.1. However, for other aspects like naturalness, content preservation and stereotypical behaviour, Llama-2 70B outperforms other models by a significant gap. One contributing factor to this gap is the comparatively lower score for localization and lower no of edits (also CSI edits) for Llama-2 70B. Since fewer items are localized, more content is likely to get preserved, fewer edits are less likely to disrupt the naturalness of the dialogue and cause stereotypical behaviour in outputs.

We verify this hypothesis based on the correlation score (using Kendall’s τ^{10}) between different aspects. Figure 7 shows that for Llama-2 70B, localization is significantly correlated to naturalness(negative) and stereotypical behaviour(positive). A strong correlation between content preservation and naturalness suggests that with content preserved, it is unlikely that the natural flow of the dialogue is compromised. However, Llama-3 70B shows no correlation between localization and naturalness indicating

¹⁰According to Botsch (2011), $|\tau| \in [0, 0.1)$ - very weak correlation, $|\tau| \in [0.1, 0.2)$ - weak correlation, $|\tau| \in [0.2, 0.3)$ - moderate correlation, and $|\tau| \in [0.3, 1.0]$ - strong correlation.

that more localized edits don’t necessarily impact naturalness, which is desirable.

Human Evaluation: LLM-based evaluation correlates well with human evaluation in all aspects. Since we are using Mixtral (Jiang et al., 2024) to automatically evaluate edits and score adapted versions of dialogs on various aspects, to justify whether an automatic evaluation is plausible, we perform human evaluation on 100 dialogs ($\approx 9\%$ of total number of dialogs to ensure statistical significance of the test) from our corpus. Mixtral¹¹ has shown remarkable performance on a number of benchmarks often outperforming closed-source LLMs like GPT-3.5¹²(Jiang et al., 2024). We take 100 dialog pairs (original and an adapted version from Llama 2 70B model) at random and ask human raters to score the adapted version given the original version on a scale of 1-5 on each aspect using the same criteria as given in the LLM prompt for dialog level aspect evaluation (Appendix B). Taking an average of scores from human annotators, we measure the correlation (Kendall’s τ) between average human rating and LLM rating. Taking inspiration from Amidei et al. (2019), we opted to use correlation rather than agreement. The agreement primarily focuses on whether two annotators exactly agree on their ratings, whereas the correlation coefficient addresses whether, “when annotator A rates an adaptation higher on an aspect, annotator B also rates that adaptation higher.” The results are shown in Table 6. For all aspects, human ratings significantly correlate with LLM ratings (all with $p\text{-value} < 0.05$), which validates the reliability of using LLM-based scoring in assessing dialog quality along these aspects for ‘India’ as the target culture.

For aspects like Naturalness and Content Preservation, this is not surprising due to the superior language understanding ability of these models, however for aspects like localisation, identifying stereotypes and offensiveness, a strong correlation can be attributed to the model’s knowledge of *Indian culture* along with the specific instructions in the prompt. However, for target cultures with lower representation in NLP, better (culturally well-informed) models need to be used

¹¹<https://huggingface.co/mistralai/Mixtral-8x7B-Instruct-v0.1>

¹²GPT-3.5-Turbo-0125




Figure 7: Correlation between aspects

in order to scale this evaluation and achieve better performance at this task.

Aspect	kendall's τ
Naturalness	0.63
Localisation	0.60
Content Preservation	0.39
Stereotypical	0.47
Offensiveness	1.00

Table 6: Correlation between Human and LLM dialog level scores

RQ 4: What insights does this offer into their parametric cross-cultural knowledge?

Through these results, it can be observed that LLMs can localize different CSI in a cross-cultural setting for the case of “USA to India” adaptation, although, the quality of content may be compromised. In many cases, efforts of localization compromise naturalness and content preservation which is not desired, and can introduce generalizations or stereotypes about the target culture. Models getting high localisation scores without much impact on other aspects like naturalness, stereotypical behaviour and content preservation indicate that the quality of localised edits is better i.e the edits are less stereotypical/offensive and they fit well in the context of the dialog, without changing the original intent or disrupting the flow of the dialog. The quality of localised edits is indicative of whether the model truly understands the technical aspects of a culture or just has a superficial knowledge of terms and concepts without much idea of how they can be used in cross-cultural scenarios such as this task of cultural adaptation.

8 Conclusion

In this paper, we explored the task of cultural adaptation within the realm of NLP. We defined the cultural elements likely to undergo transformation during adaptation. We curated a corpus of dialogues, annotating culture-specific elements across various categories and levels of foreignness, and defined the goals and aspects of cultural adaptation employing both edit-level analysis and broader, more contextual dialogue-level analysis for evaluation. We assess the performance of several open-source LLMs for cultural adaptation and analyse how these aspects tie together. We found that while modern language models are able to localise context to a target culture to a significant extent, they often struggle with reasoning over these cultural artefacts resulting in a lack of coherence within the context of dialogue which often leads to loss of original message. The ability of LLMs to localize text for a specific target culture provides a good starting point for adaptation experts to take ideas from and further refine and enhance.

References

- Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Ashutosh Dwivedi, Alham Fikri Aji, Jacki O’Neill, Ashutosh Modi, and Monojit Choudhury. 2024. [Towards measuring and modeling "culture" in llms: A survey.](#)
- Jacopo Amidei, Paul Piwek, and Alistair Willis. 2019. [Agreement is overrated: A plea for correlation to assess human evaluation reliability.](#) In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 344–354, Tokyo, Japan. Association for Computational Linguistics.
- R Botsch. 2011. Chapter 12: Significance and measures of association. *Scopes and Methods of Political Science.*
- Yong Cao, Yova Kementchedjhieva, Ruixiang Cui, Antonia Karamolegkou, Li Zhou, Megan Dare, Lucia Donatelli, and Daniel Hershcovich. 2024. Cultural adaptation of recipes. *Transactions of the Association for Computational Linguistics*, 12.
- Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Hershcovich. 2023. [Assessing cross-cultural alignment between ChatGPT and human societies: An empirical study.](#) In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 53–67, Dubrovnik, Croatia. Association for Computational Linguistics.
- Eirlys E. Davies. 2003. [A goblin or a dirty nose? The Translator](#), 9(1):65–100.

- Ryan Drapeau, Lydia B. Chilton, Jonathan Bragg, and Daniel S. Weld. 2016. [Microtalk: Using argumentation to improve crowdsourcing accuracy](#). In *AAAI Conference on Human Computation & Crowdsourcing*.
- Esin Durmus, Karina Nyugen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli. 2023. [Towards measuring the representation of subjective global opinions in language models](#).
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024. [Mixtral of experts](#).
- Ming Jiang and Mansi Joshi. 2023. [Cpopqa: Ranking cultural concept popularity by llms](#).
- David Katan. 2014. *Translating Cultures: An Introduction for Translators, Interpreters and Mediators*. Routledge.
- Reem I. Masoud, Ziquan Liu, Martin Ferianc, Philip Treleaven, and Miguel Rodrigues. 2023. [Cultural alignment in large language models: An explanatory analysis based on hofstede’s cultural dimensions](#).
- Remi Mir, Bjarke Felbo, Nick Obradovich, and Iyad Rahwan. 2019. [Evaluating style transfer for text](#). 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 495–504, Minneapolis, Minnesota. Association for Computational Linguistics.
- Joaquin Navajas, Tamara Niella, Gerry Garbulsky, Bahador Bahrami, and Mariano Sigman. 2017. [Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds](#).
- Peter Newmark. 1988. *A textbook of translation*, volume 66. Prentice hall New York.
- Denis Peskov, Viktor Hangya, Jordan Boyd-Graber, and Alexander Fraser. 2021. [Adapting entities across languages and cultures](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3725–3750, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mark Polizzotti. 2018. *Sympathy for the Traitor: A Translation Manifesto*. The MIT Press.
- Aida Ramezani and Yang Xu. 2023. [Knowledge of cultural moral norms in large language models](#).
- G. Ritzer. 1996. *The McDonaldization of Society: An Investigation Into the Changing Character of Contemporary Social Life*. Pine Forge press titles of related interest. Pine Forge Press.
- Mike Schaekermann, Joslin Goh, Kate Larson, and Edith Law. 2018. [Resolvable vs. irresolvable disagreement: A study on worker deliberation in crowd work](#). *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrubti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- H. Vermeer. 1989. [Skopos and commission in translational action](#).
- Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen tse Huang, Zhaopeng Tu, and Michael R. Lyu. 2024. [Not all countries celebrate thanksgiving: On the cultural dominance in large language models](#).
- Binwei Yao, Ming Jiang, Diyi Yang, and Junjie Hu. 2023. [Empowering llm-based machine translation with cultural awareness](#).
- Zhonghe Zhang, Xiaoyu He, Vivek Iyer, and Alexandra Birch. 2024. [Cultural adaptation of menus: A fine-grained approach](#).

Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2023. [Multilingual machine translation with large language models: Empirical results and analysis.](#)

A Limitations

English as a medium We acknowledge the fact that language strongly reflects culture. Our selection of English (for Intralingual adaptation) enabled us to focus on identifying culture-related modifications in adaptation without the complexities of translation.

Country as a proxy for culture In this study, we have selected "nation" as a proxy for culture as a proof of concept. While this choice is often made for addressing a broader national audience in such applications, it inevitably emphasizes popular aspects of culture while potentially neglecting local subcultures.

Prompt Analysis Our analysis of prompts is not exhaustive. This is due to evaluation limits as we go deeper down the levels of culture, where culture becomes less technical and more abstract as discussed in Section 4.

Single Source-Target Culture pair Our study is confined to a single source-target culture pair. While we hope to extend our study, it requires CSI annotations from people belonging to the target culture.

Evaluation on State-of-the-art LLMs We did not evaluate on state-of-the-art closed source models like GPT-3.5 and GPT-4. While comparing models is not the main goal of this study, due to our budgetary limitations as well as our commitment to open science, we decided not to evaluate on these models.

Human Evaluation Another limitation is limited human evaluation. While we have shown a correlation between human and LLM judgements on various aspects of evaluation, we still believe there is no substitute for human evaluation. However, the associated costs make large-scale studies across different cultures prohibitively expensive and unscalable.

Extent of Localisation For this study, we measured the extent to which LLMs can adapt cultural items, however, in many applications, not all CSI need to be adapted. The selective adaptation approach allows for a balance between preserving cultural authenticity and ensuring relevance and comprehension within new or diverse cultural settings.

B Prompts for LLM Evaluation

The prompt used to extract edits(at an utterance level) is given in Table 7. Using this prompt, we

can find edits corresponding to all edited CSI along with the rest of the significant edits.

The prompt used for finding translation strategy for a given CSI edit is given in Table 8.

The prompt used for scoring edits is given in Table 9.

The prompt used for scoring adapted dialog given the original dialog is given in Table 10.

C Examples used in the prompt for obtaining adaptation

The examples used in the adaptation prompt as described in Table 2 are given in Table 11.

D Example adaptations and Edits

Examples of adaptations from different models for a single dialog are given in Table 12. Table 14 shows examples of original and adapted versions of several utterances. Table 13 shows examples of Incorrect Edits found in cultural adaptations using LLM evaluation.

E Applications of Adaptation

Following are some (non-exhaustive) applications of adaptation:

Literary Translation and Entertainment Media : Literary works and Entertainment Media (Subtitles and dubbing) are adapted to maintain the original's emotional impact, and humor while replacing cultural references with equivalents that make sense to the target audience.

Advertising or Marketing : Multinational companies adapt their marketing materials to align with local values and consumer behaviour.

Training and Education Materials : Corporate training materials are often adapted to suit the cultural context of international employees. Even educational materials like storybooks are adapted to cater to different age groups.

Legal and Healthcare Documents : Medical documents are adapted to ensure patients understand their rights and the procedures. Legal Contracts are often tailored to comply with local laws.

Identify all occurrences of the lexically edited words or phrases in original vs modified form :

Examples:

Original text : “Joey Tribbiani: What are you talking about? ‘One woman’? That’s like saying there’s only one flavor of ice cream for you. Lemme tell you something, Ross. There’s lots of flavors out there. There’s Rocky Road, and Cookie Dough, and Bing! Cherry Vanilla. You could get ‘em with Jimmies, or nuts, or whipped cream! This is the best thing that ever happened to you! You got married, you were, like, what, eight? Welcome back to the world! Grab a spoon!”

Modified text : “Joey Tribbiani: What are you talking about? ‘One woman’? That’s like saying there’s only one flavor of biryani for you. Lemme tell you something, Ross. There’s lots of flavors out there. There’s Butter Chicken, and Paneer Tikka, and Paan! You could get ‘em with Naan, or rice, or raita! This is the best thing that ever happened to you! You got married, you were, like, what, eight? Welcome back to the world! Grab a spoon!”

Edits:

ice cream → biryani
Rocky Road → Butter Chicken
Cookie Dough → Paneer Tikka
Bing! Cherry Vanilla → Paan
Jimmies → Naan
nuts → rice
whipped cream → raita

Original text : “Emily: Yes, I went there due to the crowd at the vegan cafe in the arts district.”

Modified text : “Emily: Yes, I went there due to the crowd at the chai stall near the temple.”

Edits:

vegan cafe → chai stall
in the arts district → near the temple

Original text : “Rason: Want to relax by the nude beach?”

Modified text : “Rason: Want to relax by the beach and do yoga?”

Edits:

nude → # deletion
→ and do yoga # addition

Original text : “Joey: What’s the matter with you?”

Modified text : “Joey: What’s the matter with you?”

Edits:

No edit found.

Extract edits for following :

{Original utterance}
{Adapted utterance}

Table 7: Prompt for extracting relevant edits

You are a translator performing an adaptation from a foreign culture to Indian culture. Given an original dialog from a show called 'Friends' and an intralingual adapted version for the Indian audience, your task is to determine which translation strategy is used in the given edit in the context of adapted version.

In the translation of Culture-specific items, Davies defines the following translation strategies:

1. Addition is when more information is added simultaneously with the transfer from source culture to target culture, for example: *eating at Wendy's* → *eating at Wendy's, an American international fast food restaurant chain*

2. Omission is a strategy when a word or a phrase is omitted from the target culture when no equivalents can be found, for example: *getting a taco from taco bell* → *getting a taco*

3. Globalization is a strategy of exchanging cultural elements of the text with more general and neutral words, to match it with the target language culture, for example: *Kimono* → *Traditional garment*; *Hamburger* → *Burger*; *Greek yoghurt* → *Curd* etc.

4. Localization is trying to find an appropriate equivalent of the CSI in the target language, for example, *sausage* → *kebab*; *mentos* → *paan*; etc.

5. Transformation is an alteration of a CSI to another CSI which is not a local equivalent but an altered/distorted version, familiar to the target language audience, for example: *football game* → *Local cricket match*; *mentos* → *namkeen* (alteration of CSI); *pastry* → *halwa* (no close equivalent so altered the CSI); etc.

Original Dialog:

{o}

Adapted Dialog:

{a}

Based on the original dialog, adapted dialog and the given translation strategies, the strategy used in the following edit based on its context in the adapted dialog is:

{edit}

Table 8: Prompt for finding translation strategy used for a given edit

You are from India. Given an original dialog from a show called 'Friends' and an adapted version for Indian audience, your task is to **rate the given edit in the context of adapted version based on the given aspects** :

Scores for each aspect is either 0 or 1 except localisation (score can be 0,1 or 2), representing either that aspect is absent(0) or present(1) in the adapted version.

Aspects:

1. correctness (0 or 1): This aspect evaluates whether the edits made in the adapted version make sense in the context of the conversation. A score of 1 indicates that the edit seamlessly integrates into the dialogue, maintaining its natural flow, practicality and coherence. A score of 0 suggests that the edit feels out of place, forced, or does not fit well in the context.

2. localisation (0, 1, 2): This aspect assesses the degree to which the adapted version is tailored for the Indian audience. A score of 0 implies that the edit lacks any localization efforts and feels completely foreign or unrelated to the Indian context. A score of 1 signifies that the changes reflect some level of localization but are more of a globalized edit that resonates with the global audience beyond just the Indian context. A score of 2 indicates that the edit is highly localized, seamlessly integrating elements specific to the Indian culture.

3. offensiveness (0 or 1): This aspect determines whether the adapted version contains any offensive content(i.e score: 1) that is clearly inappropriate or disrespectful within the Indian cultural framework, such as mockery of Indian traditions or values, propagation of harmful stereotypes or depiction of sensitive topics without appropriate cultural sensitivity. A score of 0 indicates that no offensive content is present, while, an offensiveness score of 1 indicates the presence of offensive content that should be addressed or revised.

Original Dialog:

{}

Adapted Dialog:

{}

Based on the original dialog, adapted dialog and the aspects, please give the rating for the following edit based on its context in the adapted dialog along each aspect in Python dictionary format (no explanation) 'aspect1': scorevalue1,...:

{Edit}

Table 9: Prompt for scoring edits on different aspects.

You are from India. Given an original dialog from a show called 'Friends' and an adapted version for the Indian audience, **your task is to rate the adapted version along the following aspects**, scores for each aspect ranges from 1 to 5.

Aspects:

- 1. naturalness(1-5):** This aspect assesses the fluidity and coherence of the adapted dialogue within the context of the conversation. It evaluates whether the adapted version feels natural and seamlessly integrates cultural modifications without appearing forced or disrupting the flow of the conversation.
- 2. localisation(1-5):** This aspect assesses the degree to which the adapted version is localized/adapted for the Indian audience. A high score signifies that the changes reflect cultural nuances, language preferences, and social norms relevant to the Indian setting. A lower score means more presence of foreign items or the edits are done in a globalised/generalised way.
- 3. offensiveness(1-5):** This aspect determines whether the adapted version contains offensive content that is clearly inappropriate or disrespectful within the Indian cultural framework, such as mockery of Indian traditions or values, propagation of harmful stereotypes or depiction of sensitive topics without appropriate cultural sensitivity. The offensiveness score ranges from 1 (least offensive) to 5 (most offensive).
- 4. stereotypical(1-5):** This aspect evaluates the presence of stereotypical representations of Indian culture or its people. A high score indicates a high degree of stereotypical content, while a low score suggests a more accurate and nuanced portrayal. Some common stereotypes about India and its people include - "All Indians are vegetarians.", "All Indian women wear sarees, and all Indian men wear turbans.", "All Indian cuisine is extremely spicy and hot.", "All Indians live in extreme poverty and deprivation.", "All Indians are obsessed with Bollywood movies and stars.", "All Indians work in call centres or IT support.", etc.
- 5. content preservation (1-5):** This aspect gauges how effectively the essence, message, and emotional undertones of the original dialogue are maintained in the adapted version, irrespective of cultural style. A higher score reflects a faithful preservation of the original dialogue's message, humour, and emotional impact, while a lower score suggests significant alterations that may compromise the original intent or emotional resonance.

Scoring Criteria:

naturalness (1-5):

- 5: The adapted dialog flows seamlessly, maintaining a natural and coherent conversation.
- 4: The adapted dialog feels mostly natural and contextually fitting, with minor deviations.
- 3: Some parts of the adapted dialog may feel slightly forced or unnatural, affecting its coherence.
- 2: The coherence of the adapted dialog is disrupted by noticeable inconsistencies or unnatural phrasing.
- 1: The adapted dialog appears completely out of place or nonsensical, lacking coherence and natural flow.

localisation (1-5):

- 5: The adaptation demonstrates a deep understanding of Indian cultural nuances, language, and social norms, resulting in a highly localized portrayal.
- 4: Cultural elements are effectively integrated, enriching the adaptation with a blend of local flavour.
- 3: Localization attempts are noticeable, albeit with some inconsistencies in maintaining cultural authenticity.
- 2: Minimal localization efforts are evident, appearing globally adapted rather than localized.
- 1: The adaptation lacks any reflection of Indian culture, appearing very foreign rather than localized.

offensiveness (1-5):

- 5: Contains highly offensive material that is disrespectful and inappropriate within the Indian cultural framework.
- 4: Includes offensive content that is likely to offend or cause discomfort.
- 3: Some potentially offensive elements are present, requiring careful revision.
- 2: Minor offensive aspects that may need adjustment.
- 1: No offensiveness

stereotypical (1-5):

- 5: Contains highly stereotypical representations of Indian culture or its people.
- 4: Includes noticeable stereotypes that may perpetuate cultural biases.
- 3: Some stereotypical content is present, but not overly pervasive.
- 2: Minimal stereotypical content, with occasional instances.
- 1: No stereotypes present.

content preservation (1-5):

- 5: The adapted version faithfully retains the essence and significance of the original content.
- 4: The majority of the original content is preserved, albeit with minor adjustments.
- 3: While some changes are evident, the overall meaning remains largely intact.
- 2: Significantly altered content leads to a noticeable shift in meaning.
- 1: The original content is either lost entirely or severely distorted in the adaptation.

Original Dialog:

{}

Adapted Dialog:

{}

Based on the original dialog and the adapted dialog, please rate the adapted dialog, and give a score along each aspect with an explanation only in a JSON format {aspect: {score:, explanation:},.....}:

Table 10: Prompt for scoring adapted dialogs on different aspects.

Original Dialog 1:

Angela: Did you see the Beatles concert last night?

Mary: No, I was catching up baseball game last night on TV.

Angela: Oh! Did you eat the meatball spaghetti I made ?

Rosy: Totally! I also added some oregano and rosemary to it.

Mary: Ohkay Angela tell me, what should I wear for the date, is this skirt good?

Angela: Nope, wear the gown I gave you on last Thanksgiving.

Rosy: Yeah totally wear that. That was beautiful.

Angela: And where are you going for the date?

Mary: A nice restaurant near the White House.

Angela: Bring me gelato.

Rosy: Bye Mary!

Mary: Bye! Wish me luck, Hope I score tonight!

TRANSCRIPT NOTE: (Mary and her date meet and greet each other with a kiss)

Adapted Version 1:

Angela: Did you see Shreya Ghoshal's concert last night?

Mary: No, I was catching up cricket game last night on TV.

Angela: Oh! Did you eat the sevai I made? Rosy: Totally! I also added some gunpowder and coriander to it.

Mary: Ohkay Angela tell me, what should I wear for the date, is this kurta good?

Angela: Nope, wear the kurta I gave you on Diwali last time.

Rosy: Yeah totally wear that. That was beautiful.

Angela: And where are you going for the date?

Mary: A nice restaurant near the Red Fort.

Angela: Bring me kulfi.

Rosy: Bye Mary!

Mary: Bye! Wish me luck, Hope it goes well!

TRANSCRIPT NOTE: (Mary and her date meet and greet each other with a handshake)

Original Dialog 2:

Mark: Have you been to the new Italian restaurant downtown?

Emily: Yes, I went there due to the crowd at the vegan cafe in the arts district.

Mark: Oh! Did you try their tiramisu?

Emily: Yes, it was delicious! Nice touch of coco powder to it.

Mark: Good! Emily, I have been thinking about applying for the post of editor for Harvard Business Review.

Emily: Great Mark! Good luck, you totally deserve it.

Adapted Version 2:

Mark: Have you been to the new Kerala restaurant in the market?

Emily: Yes, I went there due to the crowd at the chai stall near the temple.

Mark: Oh! Did you try their Rava Kesari? Emily: Yes, it was delicious! Nice touch of cardamom to it.

Mark: Good! Emily, I have been thinking about applying for the post of editor for The Times of India.

Emily: Great Mark! Good luck, you totally deserve it.

Table 11: 2-shot example used in adaptation prompt

Original Dialog	Llama-2 70B	Llama-3 8B	Llama-3 70B
Franny: Hey, Monica! Monica Geller: Hey Frannie, welcome back! How was Florida? Franny: You had sex, didn't you? Monica Geller: How do you do that? Franny: Oh, I hate you, I'm pushing my Aunt Roz through Parrot Jungle and you're having sex! So? Who? Monica Geller: You know Paul? Franny: Paul the Wine Guy? Oh yeah, I know Paul. Monica Geller: You mean you know Paul like I know Paul? Franny: Are you kidding? I take credit for Paul. Y'know before me, there was no snap in his turtle for two years.	Franny: Hey, Monica! Monica Geller: Hey Frannie, welcome back! How was Goa? Franny: You had sex, didn't you? Monica Geller: How do you do that? Franny: Oh, I hate you, I'm pushing my Aunt Roz through Anjuna Flea Market and you're having sex! So? Who? Monica Geller: You know Rohan? Franny: Rohan the Yoga Guy? Oh yeah, I know Rohan. Monica Geller: You mean you know Rohan like I know Rohan? Franny: Are you kidding? I take credit for Rohan. Y'know before me, there was no spark in his chakra for two years.	Riya: Hey, Nalini! Nalini Goyal: Hey Riya, welcome back! How was Goa? Riya: You got married, didn't you? Nalini Goyal: How do you do that? Riya: Oh, I hate you, I'm pushing my Chachi through the beach and you're getting married! So? Who's the lucky guy? Nalini Goyal: You know Rohan? Riya: Rohan the IT guy? Oh yeah, I know Rohan. Nalini Goyal: You mean you know Rohan like I know Rohan? Riya: Are you kidding? I take credit for Rohan. Y'know before me, there was no spark in his eyes for two years.	Franny: Hey, Monica! Monica Geller: Hey Frannie, welcome back! How was Goa? Franny: You had a romantic getaway, didn't you? Monica Geller: How do you do that? Franny: Oh, I hate you, I'm pushing my Aunt Roz through the spice plantation and you're having a romantic time! So? Who? Monica Geller: You know Rohan? Franny: Rohan the Cafe Owner? Oh yeah, I know Rohan. Monica Geller: You mean you know Rohan like I know Rohan? Franny: Are you kidding? I take credit for Rohan. Y'know before me, there was no spark in his life for two years.

Table 12: Example of adaptations from different models for a single dialog

Original	Adaptation	Reason (LLM Verbal Rationale)
... Paul the Wine Guy: Well, you might try accidentally breaking something valuable of hers, say her- Monica Geller: -leg ? Paul the Wine Guy: That's one way! Me, I- I went for the watch. Paul the Wine Guy: Well, you might try accidentally breaking something valuable of hers, say her- Monica Geller: -Bindi ? Paul the Wine Guy: That's one way! Me, I- I went for the watch. ... # (Llama 2 70B)	A bindi is typically a decorative mark or dot applied to the forehead, commonly used in South Asian cultures. It is usually made from a small adhesive sticker, kumkum (vermilion), or other cosmetic materials. Technically, a bindi itself cannot be "broken". It might be peeled off or damaged, but not broken.
... Paul the Wine Guy: Ever since she walked out on me, I, uh... Monica Geller: What?..... What, you wanna spell it out with noodles ? Raj: Ever since she walked out on me, I, uh... Monica: What?.. What, you want to spell it out with naan ? ... # (Llama 2 70B, Llama 3 8B) ... Raj: Ever since she walked out on me, I, uh... Monica: What?.. What, you want to spell it out with gulab jaamun ? ... # (Llama 3 70B)	The original phrase "spell it out with noodles" is a humorous expression that makes sense as noodles might be used to spell out words as a playful or whimsical act. "Naan" is a type of bread commonly eaten in South Asian cuisine. Spelling out something with naan is not common or intuitive, making the adaptation feel forced. Gulab jamun is a popular Indian dessert, but it is a round, syrup-soaked sweet that doesn't lend itself to being used to spell out words. This makes the phrase less intuitive and relatable.

Table 13: Few examples of Incorrect Edits (in red) found in cultural adaptations by LLM evaluator. These mistakes often arise due to a lack of understanding and reasoning about cultural objects.

Original Lines	Adapted Lines
<p>...</p> <p>Chandler Bing: I can't believe you would actually say that. I would much rather be Mr. Peanut than Mr. Salty.</p> <p>Joey Tribbiani: No way! Mr.Salty is a sailor, all right, he's got to be, like, the toughest snack there is.</p> <p>Ross Geller: I don't know, you don't wanna mess with corn nuts. They're craazy.</p> <p>...</p>	<p>...</p> <p>Chandler Bing: I can't believe you would actually say that. I would much rather be Haldiram's than Bikano.</p> <p>Joey Tribbiani: No way! Bikano is a strong brand, all right, they've got to be, like, the toughest snack there is.</p> <p>Ross Geller: I don't know, you don't wanna mess with murmure. They're craazy.</p> <p>... # (Llama 3 70B)</p>
<p>...</p> <p>Chandler Bing: Dehydrated Japanese noodles under fluorescent lights... does it get better than this?</p> <p>...</p>	<p>...</p> <p>Chandler Bing: Maggi under the streetlights... does it get better than this?</p> <p>... # (Llama 3 8B)</p>
<p>...</p> <p>Monica Geller: Are these the shoes?</p> <p>Rachel Green: Yes. Paolo sent them from Italy.</p> <p>...</p>	<p>...</p> <p>Monica Geller: Are these the sandals?</p> <p>Rachel Green: Yes. Rohan sent them from Jaipur.</p> <p>... # (Llama 3 8B)</p>
<p>...</p> <p>Joanna: Oh. And Rachel has been really incredible in getting my morning bagel for me. It's amazing how she gets it right almost every time! ...</p>	<p>...</p> <p>Joanna: Oh. And Rachel has been really incredible in getting my morning chai for me. It's amazing how she gets it right almost every time!</p> <p>... # (Llama 2 70B)</p>
<p>...</p> <p>Rachel Green: Alright, c'mon! Let's play Twister!</p> <p>...</p>	<p>...</p> <p>Rachel Green: Alright, c'mon! Let's play Antakshari!</p> <p>... # (Llama 2 70B)</p>
<p>...</p> <p>Joey Tribbiani: Well, Chandler's old roommate was Jewish, and these are the only candles we have, so... Happy Chanukah, everyone. ...</p>	<p>...</p> <p>Joey Tribbiani: Well, Chandler's old roommate was Hindu, and these are the only diyas we have, so... Happy Diwali, everyone.</p> <p>... # (Llama 3 70B)</p>
<p>...</p> <p>Monica Geller: He is so cute. So, where did you guys grow up?</p> <p>Angela Delveccio: Brooklyn Heights.</p> <p>Bob: Cleveland.</p> <p>...</p>	<p>...</p> <p>Monica Geller: He is so cute. So, where did you guys grow up?</p> <p>Angela Delveccio: Bandra.</p> <p>Bob: Ahmedabad.</p> <p>... # (Llama 3 70B)</p>

Table 14: Few examples of original and adapted versions of several utterances in the corpus of dialogs. Major edits are highlighted.