

On the effective transfer of knowledge from English to Hindi Wikipedia

Paramita Das
IIT Kharagpur
dasparamita1708@gmail.com

Amartya Roy
Bosch
Amartya.Roy@in.bosch.com

Ritabrata Chakraborty
Osmania University
ritabratac1997@gmail.com

Animesh Mukherjee
IIT Kharagpur
animeshm@cse.iitkgp.ac.in

Abstract

Although Wikipedia is the largest multilingual encyclopedia, it remains inherently incomplete. There is a significant disparity in the quality of content between high-resource languages (HRLs, e.g., English) and low-resource languages (LRLs, e.g., Hindi), with many LRL articles lacking adequate information. To bridge these content gaps we propose a lightweight framework to enhance knowledge equity between English and Hindi. In case the English Wikipedia page is not up-to-date, our framework extracts relevant information from external resources readily available (such as English books), and adapts it to align with Wikipedia’s distinctive style, including its *neutral point of view* (NPOV) policy, using in-context learning capabilities of large language models. The adapted content is then machine-translated into Hindi for integration into the corresponding Wikipedia articles. On the other hand, if the English version is comprehensive and up-to-date the framework directly transfers knowledge from English to Hindi. Our framework effectively generates new content for Hindi Wikipedia sections, enhancing Hindi Wikipedia articles respectively by 65% and 62% according to automatic and human judgment-based evaluations.

1 Introduction

Despite the wide usage of the multilingual content of Wikipedia, significant knowledge gaps exist across different language editions of Wikipedia (Miquel-Ribé and Laniado, 2018), creating an information divide. For instance, the Hindi Wikipedia, with only 163 thousand articles as of 2023, contrasts sharply with the English Wikipedia’s 6.8 million articles¹, despite Hindi being the third most spoken language globally. It is evident that in many low-resource languages

(LRLs), Wikipedia often lacks pages on important global topics or notable individuals due to the limited participation of community editors. This disparity limits the engagement of LRL communities with online resources and educational content. Moreover, Wikipedia articles on the same topic often differ significantly across languages (Miquel-Ribé and Laniado, 2020; Roy et al., 2020) due to factors such as cultural relevance and the varied expertise of contributors. Addressing these disparities is crucial for achieving *knowledge equity*, a concept emphasized by the Wikimedia Foundation (Redi et al., 2020). Existing studies (Zhang et al., 2024; Shao et al., 2024) primarily focus on generating full-length Wikipedia articles in English, which restricts research efforts for LRLs. Recent studies (Taunk et al., 2023; Shao et al., 2024; Maurya and Desarkar, 2023) have made significant progress in automated methods for cross-lingual knowledge transfer, particularly in generating full-length articles in LRLs. However, many of the existing approaches focus on generating articles from scratch, which is often less effective for enriching existing articles and overlooks the collaborative nature of knowledge creation. On large crowd-sourced platforms, such as Wikipedia, collaborative efforts, especially human-curated content hold greater importance than automatically generated information. Often, Wikipedia articles on specific topics in LRLs require enrichment in certain sections compared to their counterparts in high-resource languages (HRLs). To the best of our knowledge, no prior work has specifically addressed the enrichment of section-specific content of Wikipedia articles in LRLs. To address this issue, we propose a lightweight approach that enriches section-specific content while preserving existing human-authored material and augmenting it with carefully integrated, automatically generated knowledge using standard NLP techniques. Our framework utilizes the standard retrieval aug-

¹https://en.wikipedia.org/wiki/List_of_Wikipedias

mented generation (RAG) framework to extract relevant information from a web corpus for a given section title, followed by machine translation of the extracted content into LRL. Additionally, the framework addresses additional challenges, especially finding suitable references relevant to the section. In our experiment, we consider English as a representative HRL and Hindi as a representative LRL. The key research question drives our experiment: How can we automatically transfer knowledge from the more enriched language version (HRL) to the less enriched one (LRL) in Wikipedia, given a Wikipedia article on a topic t ? Our contributions are as follows.

- We propose a multistage approach to extract knowledge from biographical writings about popular figures, transform this text to adhere to the NPOV guidelines and incorporate it into low-resource Wikipedia articles, e.g., Hindi. This approach uses the *WikiTransfer* framework to identify and transfer relevant content from English Wikipedia to Hindi Wikipedia, leveraging IndicTrans².
- We manually curate 103 biographical writings relevant to corresponding Wikipedia articles as external knowledge sources, which can serve as a rich source of factual information.
- Our lightweight framework efficiently updates Hindi Wikipedia articles by adding coherent and new information. Our rigorous evaluation through both automated and crowd-based assessments demonstrates an improvement of 65% and 62%, respectively, in terms of integrating new information. Our code and dataset for reproducing similar content are available at <https://github.com/paramita08/wikiTransfer>.

The applicability of our framework has been demonstrated within the Wikipedia domain; however, its potential usage extends to large-scale industrial applications where enriching local knowledge repositories using open-source automated systems, such as large language models (LLMs), is infeasible due to proprietary data or domain-specific constraints. In such scenarios, our pipeline—leveraging semantic search based retrieval systems, i.e., RAG, and further adapting the system to be used in low-resource settings through domain-sensitive LLMs (Liu et al., 2023)—offers a practical and scalable solution. Our approach is

particularly relevant for industries dealing with sensitive or specialized knowledge repositories, where conventional generative AI models may fall short.

2 Related work

The research community has expanded NLP research in multilingual settings including more languages, especially non-English, and smaller low-resource language editions (Wang et al., 2023). In case of Wikipedia, many researchers have examined differences between different language editions of Wikipedia from the standpoints of content (i.e., text, image), readers (Arora et al., 2022), and editors (Bipat et al., 2018) as well. Text diversity in Wikipedia has collectively demonstrated that textual content about the same topic is highly diverse across language editions (Hecht and Gergle, 2010; Roy et al., 2020). Different language editions of Wikipedia serve very different communities (Johnson et al., 2021; Lemmerich et al., 2019) and thus often cover very different topics (Paramita et al., 2017). This information gap results in variations in the quality and quantity of content (Lewoniewski et al., 2017), presumably affecting the vocabulary and ability of language models trained on Wikipedia to handle different topics accurately. A large body of work (Adar et al., 2009; Wulczyn et al., 2016; Paramita et al., 2017) based on vanilla NLP approaches tried addressing the information asymmetry between different language editions. With the advancement of generative AI, recent works (Agarwal et al., 2020; Shivansh et al., 2023; Guo et al., 2024) have focused on content alignment and content transfer in low-resource languages from scratch. In the case of languages with limited or poor translation resources, authors in (Paramita et al., 2017) proposed a lightweight approach to measure cross-lingual similarity in Wikipedia using section headings rather than the entire Wikipedia article, and language resources derived from Wikipedia and Wiktionary to perform translation. Although existing research works are valuable, but lacks an end-to-end pipeline for transferring content from high-resource to low-resource languages, limiting efforts to bridge Wikipedia’s content gap. Our work addresses this issue.

3 Dataset description

We employ a systematic approach to collect Wikipedia articles, which are available in both English and Hindi versions. We also anchor on the

²<https://ai4bharat.iitm.ac.in/indic-trans2/>

content of external resources for a subset of articles. In this work, we utilize a dataset comprising biographies, i.e., articles of Wikipedia category *people* sourced from Wikipedia. As the biography articles in Wikipedia follow a predefined structure across multiple languages, we concentrate on biography articles of renowned persons as the dataset for our experiment. Although our experiments have focused solely on biography articles, our framework is versatile and can be readily applied to other types of Wikipedia articles, such as articles covering technical concepts or geographical locations. This extension is feasible as long as a sufficient digital corpus on the topic is available to serve as an external resource for the RAG module of our pipeline.

Collection of Wikipedia articles: Authors in (Beytía et al., 2022) published a dataset of Wikipedia biographies in the ten most widely spoken languages, including English and Hindi³. We use this dataset to extract biographies, along with their Wikidata IDs, which serve as unique identifiers across language versions. For instance, Serena Williams’ biography has the Wikidata ID *Q11459*, allowing retrieval in all available languages. First, we collected a set of 21,340 biography articles in both Hindi and English versions from the above-mentioned dataset. Using the MediaWiki API⁴, we retrieve and pre-process the current version of wikitext of these articles. Section headings are extracted using the Wikipedia Python package⁵, excluding sections like *See also*, *References*, and *External links*, and *Inline citations*.

Collection of article quality: We utilize article quality as an indicator to determine which language version (English and Hindi) contains more enriched information between the two. Therefore, we aim to gather the quality scores for each language version of Wikipedia articles. Using the dataset from (Das et al., 2024), we collect quality scores (ranging from 0 to 1) for Hindi and English Wikipedia articles. We identify a subset of 17,226 articles where Hindi scores are lower than English, serving as our candidate set for further experiments. Since no rigid quality class hierarchy exists for Hindi articles on Wikipedia, we use this language-agnostic quality score. Next, we extract quality categories (FA, A, GA, B, C, Start, Stub) for English articles using

Quality class	# of Articles	# of Biographies
FA	235	0
A	6	0
GA	485	13
B	1930	51
C	3428	38
Start	6625	0
Stub	4517	0

Table 1: Filtered dataset – articles categorized in quality classes and biographical writings extracted for the corresponding classes.

ORES. For articles in the FA category (according to English Wikipedia), we directly use their content to improve lower-quality Hindi versions. For other categories, English articles are first enhanced using external resources, followed by transferring the improved content to Hindi, thus ensuring that the highest-quality information is used to enhance Hindi articles.

Collection of external resources: We use online digital library *Archive*⁶, to source biographical writings for our enhancements. *Archive* offers a vast collection of scanned historical books, making it ideal for our needs.

Automated search: We first construct a search query using the title of each biography article to locate the corresponding page on *Archive*. For a given biography, say *P*, the query retrieves biographical writings, say *bio*. We use the requests library and the HTTP GET method to extract the page content. If the keyword ‘biography’ is found, the response is considered valid, ensuring relevant results from *Archive*.

Manual verification: Due to name ambiguity and automated search limitations, many results contain irrelevant information. To address this, a post-graduate student who frequently uses Wikipedia manually verified the collected links. This ensures the quality and relevance of the biographical writings used. We download the verified biographical writings in text format to enrich Wikipedia articles, aiming to improve the quality of both English and Hindi biographies. Thus, our curated dataset includes Wikipedia articles in both English and Hindi, their quality scores, and a set of biographical writings extracted from external sources. The dataset statistics is noted in Table 1.

4 Experimental pipeline

We propose an end-to-end pipeline to transfer knowledge from English articles to their corresponding Hindi versions. The pipeline includes

³<https://www.visualcapitalist.com/100-most-spoken-languages/>

⁴https://www.mediawiki.org/wiki/API:Get_the_contents_of_a_page

⁵<https://pypi.org/project/wikipedia/>

⁶www.archive.org

the framework WikiTransfer and additional modules for external knowledge augmentation and POV correction to enhance the Hindi version of articles.

4.1 WikiTransfer

WikiTransfer first identifies semantically similar section titles between English and Hindi Wikipedia articles. To map sections, we translate Hindi titles to English using IndicTrans, compute embeddings for all titles, and measure cosine similarity between pairs of Hindi and English titles of every article. For instance, for a Hindi title denoted as t_h , we compute the similarity with the embedding of m English section titles of an article p . For each Hindi title, the most similar English section is identified as the source for content transfer. We use the sentence transformers model all-MiniLM-L12-v2⁷ for embedding computation. Section pairs with a similarity score above a threshold of 0.44 (mean similarity) are selected as mapped sections. After matching section titles, we analyze the content of the mapped sections for coherence. We compute embeddings of the section content (Hindi and English) using multilingual e5-large⁸ and calculate cosine similarity. Section-content pairs with similarity scores above a threshold of $\mu + \sigma$ (mean: 0.89, std dev: 0.06) are considered similar.

Content augmentation: After mapping sections and content between English and Hindi Wikipedia articles, this step involves translating English content into Hindi using the IndicTrans model (Gala et al., 2023). Translated sentences are appended to the existing content in the mapped Hindi section. Before incorporating the translated sentences, we apply a two-step filtering: (1) discard short translations with one or two words to avoid errors, and (2) use the multilingual e5-base model of the sentence transformer to identify the top three semantically related sentences for each existing Hindi sentence. Likewise previous mapping scheme, a cosine similarity score is calculated for each x in the existing Hindi sentences with individual translated sentences in $Hindi(e)$, and we select the top three sentences among the sorted (in descending order) $Hindi(e)$ sentences that belong to the range of μ and $\mu + \sigma$ of the similarity scores. This way, we pick up three sentences that are somewhat dissimilar from the existing sentence x , thus avoiding

redundancy. If a sentence is selected for multiple matches, it is appended only once, creating a reduced and informative set of translated sentences to be added.

4.2 Content extraction from biography

English articles that belong to quality classes other than FA might require additional information to enhance their quality. Therefore, we first gather information from external resources of biography writings (as described in section 3) for such articles in our dataset and add them to the appropriate sections for further processing using the WikiTransfer framework. To extract information from these external narratives aligned with the content of the articles, we employ the standard RAG method. For each English article, we select the most recent biography, split the text into chunks using the *RecursiveTextSplitter*⁹ function, and embed each chunk using sentence-bert¹⁰ embeddings which are stored in the vector database-CHROMADB. Now, for the given English article E_p with m sections, we provide the content of the section t_i where $i \in 1, 2, \dots, m$ as query and external narratives W as the input to the *retriever* module of the RAG pipeline. We use maximum marginal relevance (MMR) based search to retrieve top k chunks (we fix k to 3) relevant to the query. Out of these retrieved chunks, we utilize a suitable prompt (Llama-3(8B)-Instruct model), which identifies which chunk is the most relevant to the given section content (please see the prompt in Appendix [A]). Instead of using the RAG text generator, we use a POV rectifier module (as discussed below) to refine the content.

4.3 POV correction

Alongside Wikipedia’s openness, a fundamental pillar of its success is its commitment to the NPOV policy, which ensures that facts should be presented fairly and impartially. According to this policy, Wikipedia prohibits sentences that contain perspective-specific or biased language, such as expressions of praise, criticism, or other sentiments that reflect the editor’s personal feelings. Given that we are extracting content from biographies, which may include subjective language, there is a risk that the new content could violate Wikipedia’s NPOV¹¹

⁷<https://huggingface.co/sentence-transformers/all-MiniLM-L12-v2>

⁸<https://huggingface.co/intfloat/multilingual-e5-large>

⁹<https://github.com/langchain-ai/langchain>

¹⁰<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

¹¹<https://tinyurl.com/cb7yv3tt>

standards. Therefore, to adhere to Wikipedia’s NPOV policy, we identify and remove subjective biases, named as framing bias and epistemological bias (Recasens et al., 2013) from individual sentences extracted from the biographies if they exist and rephrase them accordingly. In this study we have tried to leverage the power of LLMs for the generation of Wiki-style content. The most popular methods to use LLM in such downstream tasks are • supervised finetuning (Jiang et al., 2019) • in-context learning (Sahoo et al., 2024). We have performed our experiments with LLama-3(8B) instruct model (AI@Meta, 2024) for both these setups.

Supervised fine-tuning (SFT): For this setup, we have fine-tuned the LLama-3(8B)-instruct model using the WNC and WikiBias¹² corpus in a supervised fashion to obtain a supervised fine-tuned model.

In-context learning (ICL): For this setup, we have used off-the-shelf instruction-tuned models namely LLama-3(8B) & LLama-3(70B). We have used a generic prompt (please see the prompt in Appendix [B]) to generate a debiased sentence given a biased sentence. Specifically, we have tried • zero-shot (only the instruction) and • few-shot (Parnami and Lee, 2022) (a few examples are used to describe the task to the model) prompting for the generation of Wiki-style neutral content.

Evaluation: Both of these configurations are conducted on a sample of test data comprising 431 biased sentences and their neutral counterparts¹³ and the evaluation of generated neutral content are computed using three reference-based metrics—BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and BERTScore (Zhang et al., 2019). As evidenced in Table 2, ICL consistently outperformed SFT across all three the reference-based metrics. Hence we have used the ICL few-shot setup (5 examples have been used in the prompt) and Llama-3(70B) in generating neutral content for the extracted external book content as mentioned in the above section.

5 Results

We assess the LLM-generated neutral Wiki-style content and the machine-translated Hindi content

¹²We have used $\sim 2k$ from WikiBias (*Trainmanual*) (Zhong, 2021) and 10k biased and neutral sentence pairs randomly sampled from WNC corpus (Pryzant et al., 2020) as training data.

¹³We have utilized WikiBias test data here.

Model	Methods	BLEU	METEOR	BERT
Llama-3(8B)[SFT]	Zero-shot	0.27	0.5	0.94
	Few-shot	0.35	0.65	0.92
Llama-3(8B)[ICL]	Zero-shot	0.25	0.6	0.93
	Few-shot	0.35	0.66	0.95
Llama-3(70B)[ICL]	Zero-shot	0.24	0.57	0.93
	Few-shot	0.4	0.68	0.95

Table 2: Evaluation score of Llama-3 on test data.

through automatic metrics and human evaluation.

5.1 Automatic evaluation

To evaluate the relatedness and quality of the newly generated content with the pre-existing Hindi text, we employ the E-A-T framework proposed by (Sugandhika and Ahangama, 2022) for the information quality assessment of Wikipedia content, and three important factors of the framework are – Expertise (E), Authority (A), and Trust (T). For the purpose of assessment of the machine-translated content, we valued E more than the other two factors, which comprise the following components –

Informativeness ($Info$) = $0.12 * \text{page-size} + 0.151 * \text{\#sentences} + 0.154 * \text{\#words} + 0.155 * \text{\#complex-words}$;

Readability ($Read$) = $0.213 * \text{Flesch-Kincaid-grade-level} + 0.185 * \text{Coleman-Liau-index} + 0.26 * \text{\%complex-words} + 0.253 * \text{avg-syllables-per-word}$;

Understandability (Und) = $0.393 * \text{Gunning-Fog-score} + 0.352 * \text{SMOG-index} + 0.181 * \text{automated-readability-ndex} + 0.344 * \text{avg-words-per-sentence}$;

Finally, E is measured in terms of the Quality of a Wikipedia page content which is defined as: $Quality (Qual) = 0.255 * \text{Informativeness} + 0.654 * \text{Readability} + 0.557 * \text{Understandability}$. Informativeness represents the size of the textual content on the Wikipedia page; readability and understandability provide insights into the linguistic quality of the article content. We perform reverse translation of newly generated and existing Hindi content into English and then evaluate it using the above-mentioned approach.

Results of automatic evaluation: Due to limited resources for evaluating Hindi text quality, we assess the quality in English. We perform reverse translation of both the newly generated Hindi

Original biased sentence: Blacks never listen to their parents.		
Score	Rules	Example
1	Complete bias removal	People do not always listen to their parents.
2	Complete bias removal + Keeping the meaning (context) same	Some people never listen to their parents.
3	Complete bias removal + keeping the meaning (context) same + fluency	It is not uncommon for individuals to disregard parental advice.

Table 3: **Scoring metric.** Details of the scoring metrics used for annotation along with examples based on a biased sentence. The original biased sentence is taken from CrowS-Pairs dataset (Nangia et al., 2020).

Type	Info		Read		Und		Qual	
	c_{old}	c_{new}	c_{old}	c_{new}	c_{old}	c_{new}	c_{old}	c_{new}
FA	53.35(50.65)	114.02(73.84)	4.66(0.79)	4.90(0.62)	17.09(4.02)	18.05(2.84)	26.17(13.69)	42.39(19.04)
Non-FA	62.94(45.09)	136.17(61.7)	4.92(0.72)	5.17(0.62)	17.01(3.65)	17.74(2.75)	28.74(11.76)	47.99(15.66)

Table 4: Human evaluation on the generated machine-translated Hindi content based on three metrics.

content c_{new} and the existing Hindi content c_{old} (which is the set of existing Hindi sentences in a section) and then evaluate them using the above-mentioned approach. We compute the metrics for individual sections and average over all the sections of the articles under consideration. The scores obtained for the two groups – (1) FA and (2) non-FA (GA, B, and C quality articles together) using automatic evaluation are tabulated in Table 4. Overall, we observe that the enhanced content is superior to the old content in terms of all the metrics for both groups. Since the standard deviation obtained in the case of Informativeness is large, we provide further division of the metrics (mean and standard deviation) in Table 5 in the Appendix.

5.2 Human evaluation

Assessment of LLM-generated NPOV text: To evaluate the neutrality of the LLM-generated text, we conduct the human assessment on 50 randomly sampled sentences from our dataset, comparing the original sentence from external resources with the NPOV version generated by the Llama-3(70B) model. Two evaluators assigned scores using the scoring metric defined in Table 3. The average score ($Score_{new}$) assessed by the two evaluators are 82.85% and 77.14%, respectively, showing that the neutralized content is suitable for augmenting the target Hindi sections.

Overall assessment: We evaluate the content generated by our pipeline in two scenarios: (a) augmenting content using only FA articles, and (b) augmenting content from non-FA articles along with external sources. The evaluation focuses on three qualitative metrics – *informativeness*, *readability*, and *coherence* – each rated on a scale from 1 to 3. Informativeness, in this context, indicates the ability of a piece of text to provide useful information

and comprehensive content. Readability measures the effort required by the reader to read and understand a piece of information. If the vocabulary and sentence structure in the text are complex, the difficulty of reading increases. Coherence represents the logical flow between sentences in a text, ensuring that they naturally follow one another to form meaningful content. Seven Hindi-speaking evaluators conducted this assessment. For each metric, they compared the original Hindi content, c_{old} , consisting of existing Hindi sentences (h), with the newly generated Hindi content, c_{new} , of the corresponding section. Evaluators assigned scores based on improvement, no change, or decline, labeled as 3, 2, and 1, respectively where a higher score indicates greater improvement of c_{new} compared to c_{old} . We randomly sample 35 sections from our dataset to evaluate the content generated by our framework: 10 from the FA category and 25 from non-FA quality categories (5 GA, 10 B, and 10 C). The average informativeness, readability and coherence across the seven human judges respectively are – (FA: 2.67, non-FA: 2.68), (FA: 2.46, non-FA: 2.50), and (FA: 2.34, non-FA: 2.32)¹⁴. Thus we observe that for all the metrics and for both FA and non-FA categories, the average judgements are always 2.3+ indicating the newly generated content is reasonably good in terms of the three metrics, especially informativeness and readability. We find significant improvement in informativeness for both FA and non-FA groups, suggesting effective addition of relevant knowledge to existing sections (see Appendix [Figure 1, Figure 2, Figure 3, Figure 4] for examples). Given our multi-label evaluation scheme and involving multiple annotators, we compute inter-annotator agreement using Fleiss’ Kappa method (Fleiss and Cohen, 1973). We obtain κ values of 0.61, 0.53, and 0.54 for informativeness,

¹⁴Please see Table 6 in the Appendix for ratings obtained for each of the seven individual annotators.

readability, and coherence, respectively, indicating a moderate to substantial level of agreement among annotators in assessing the generated content.

6 Conclusion

In this paper, we presented a lightweight framework to produce content that is substantially superior to the existing content for Hindi articles. From FA-quality English articles, we directly translated relevant content to their corresponding Hindi counterparts. For non-FA articles, we first extracted relevant content from external sources and adapted these to Wikipedia’s NPOV style using the in-context learning capabilities of LLMs. Finally, the combined knowledge (existing and newly extracted content) in English is machine-translated into Hindi. We performed a comprehensive evaluation based on automated metrics and human assessments to demonstrate that the added content is informative, readable, and coherent. Our proposed pipeline is adaptable to any combination of HRL and LRL pairs. While the automated approach helps bridge information gaps in low-resource languages, it may risk overshadowing subtle cultural elements. To mitigate this, language-specific domain experts should perform thorough manual reviews before content integration.

7 Limitations

Despite the promising results, our study has certain limitations. Our manual verification process, while crucial for ensuring content quality, is inherently subjective and may result in inconsistencies in evaluating relevance and accuracy. Furthermore, although the dataset of personal narratives is diverse, it may not fully represent all lesser-known biographies, which could limit the generalizability of our approach. Future research should aim to integrate more diverse sources and develop automated verification techniques to address these limitations.

8 Ethical consideration

The biographical writings used for data collection were sourced from publicly accessible digital libraries, adhering to copyright regulations and respecting intellectual property rights. All human annotators involved in the manual verification process participated voluntarily. No personally identifiable information was gathered from the annotators, ensuring their privacy and anonymity. Additionally,

we took extensive measures to prevent the inclusion of sensitive or potentially harmful content in the enhanced Wikipedia articles.

References

- Eytan Adar, Michael Skinner, and Daniel S Weld. 2009. Information arbitrage across multi-lingual wikipedia. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 94–103.
- Pushkal Agarwal, Miriam Redi, Nishanth Sastry, Edward Wood, and Andrew Blick. 2020. Wikipedia and westminster: Quality and dynamics of wikipedia pages about uk politicians. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, pages 161–166.
- AI@Meta. 2024. [Llama 3 model card](#).
- Akhil Arora, Martin Gerlach, Tiziano Piccardi, Alberto García-Durán, and Robert West. 2022. Wikipedia reader navigation: When synthetic data is enough. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, pages 16–26.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Pablo Beytía, Pushkal Agarwal, Miriam Redi, and Vivek K. Singh. 2022. [Visual gender biases in wikipedia: A systematic evaluation across the ten most spoken languages](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1):43–54.
- Taryn Bipat, David W McDonald, and Mark Zachry. 2018. Do we all talk before we type? understanding collaboration in wikipedia language editions. In *Proceedings of the 14th International Symposium on Open Collaboration*, pages 1–11.
- Paramita Das, Isaac Johnson, Diego Saez-Trumper, and Pablo Aragón. 2024. Language-agnostic modeling of wikipedia articles for content quality assessment across languages. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 1924–1934.
- Joseph L Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and psychological measurement*, 33(3):613–619.
- Jay Gala, Pranjal A Chitale, Raghavan AK, Varun Gumma, Sumanth Doddapaneni, Aswanth Kumar, Janki Nawale, Anupama Sujatha, Ratish Puduppully, Vivek Raghavan, et al. 2023. Indictrans2: Towards

- high-quality and accessible machine translation models for all 22 scheduled indian languages. *arXiv preprint arXiv:2305.16307*.
- Ping Guo, Yubing Ren, Yue Hu, Yunpeng Li, Jiarui Zhang, Xingsheng Zhang, and He-Yan Huang. 2024. Teaching large language models to translate on low-resource languages with textbook prompting. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15685–15697.
- Brent Hecht and Darren Gergle. 2010. The tower of babel meets web 2.0: user-generated content and its applications in a multilingual context. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 291–300.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019. Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. *arXiv preprint arXiv:1911.03437*.
- Isaac Johnson, Florian Lemmerich, Diego Sáez-Trumper, Robert West, Markus Strohmaier, and Leila Zia. 2021. Global gender differences in wikipedia readership. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, pages 254–265.
- Florian Lemmerich, Diego Sáez-Trumper, Robert West, and Leila Zia. 2019. Why the world reads wikipedia: Beyond english speakers. In *Proceedings of the twelfth ACM international conference on web search and data mining*, pages 618–626.
- Włodzimierz Lewoniewski, Krzysztof Węcel, and Witold Abramowicz. 2017. Relative quality and popularity evaluation of multilingual wikipedia articles. In *Informatics*, volume 4, page 43. MDPI.
- Mingjie Liu, Teodor-Dumitru Ene, Robert Kirby, Chris Cheng, Nathaniel Pinckney, Rongjian Liang, Jonah Alben, Himyanshu Anand, Sanmitra Banerjee, Ismet Bayraktaroglu, et al. 2023. Chipnemo: Domain-adapted llms for chip design. *arXiv preprint arXiv:2311.00176*.
- Kaushal Maurya and Maunendra Desarkar. 2023. Towards low-resource language generation with limited supervision. In *Proceedings of the Big Picture Workshop*, pages 80–92.
- Marc Miquel-Ribé and David Laniado. 2018. Wikipedia culture gap: quantifying content imbalances across 40 language editions. *Frontiers in physics*, 6:54.
- Marc Miquel-Ribé and David Laniado. 2020. The wikipedia diversity observatory: A project to identify and bridge content gaps in wikipedia. In *Proceedings of the 16th International Symposium on Open Collaboration*, pages 1–4.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R Bowman. 2020. Crows-pairs: A challenge dataset for measuring social biases in masked language models. *arXiv preprint arXiv:2010.00133*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Monica Lestari Paramita, Paul Clough, and Robert Gaizauskas. 2017. Using section headings to compute cross-lingual similarity of wikipedia articles. In *Advances in Information Retrieval: 39th European Conference on IR Research, ECIR 2017, Aberdeen, UK, April 8-13, 2017, Proceedings 39*, pages 633–639. Springer.
- Archit Parnami and Minwoo Lee. 2022. Learning from few examples: A summary of approaches to few-shot learning. *arXiv preprint arXiv:2203.04291*.
- Reid Pryzant, Richard Diehl Martinez, Nathan Dass, Sadao Kurohashi, Dan Jurafsky, and Diyi Yang. 2020. Automatically neutralizing subjective bias in text. In *Proceedings of the aaai conference on artificial intelligence*, volume 34, pages 480–489.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st annual meeting of the Association for Computational Linguistics (volume 1: long papers)*, pages 1650–1659.
- Miriam Redi, Martin Gerlach, Isaac Johnson, Jonathan Morgan, and Leila Zia. 2020. A taxonomy of knowledge gaps for wikimedia projects (second draft). *arXiv preprint arXiv:2008.12314*.
- Dwaipayana Roy, Sumit Bhatia, and Prateek Jain. 2020. A topic-aligned multilingual corpus of wikipedia articles for studying information asymmetry in low resource languages. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2373–2380.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.
- Yijia Shao, Yucheng Jiang, Theodore Kanell, Peter Xu, Omar Khattab, and Monica Lam. 2024. Assisting in writing wikipedia-like articles from scratch with large language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6252–6278.
- Subramanian Shivansh, Maity Ankita, Jain Aakash, Singh Bhavyajeet, Gupta Harshit, Khanna Lakshya,

- and Varma Vasudeva. 2023. Cross-lingual fact checking: Automated extraction and verification of information from wikipedia using references. In *Proceedings of the 20th International Conference on Natural Language Processing (ICON)*, pages 828–831.
- Chinthani Sugandhika and Supunmali Ahangama. 2022. Assessing information quality of wikipedia articles through google’s eat model. *IEEE Access*, 10:52196–52209.
- Dhaval Taunk, Shivprasad Sagare, Anupam Patil, Shivansh Subramanian, Manish Gupta, and Vasudeva Varma. 2023. Xwikigen: Cross-lingual summarization for encyclopedic text generation in low resource languages. In *Proceedings of the ACM Web Conference 2023*, pages 1703–1713.
- Wenxuan Wang, Zhaopeng Tu, Chang Chen, Youliang Yuan, Jen-tse Huang, Wenxiang Jiao, and Michael R Lyu. 2023. All languages matter: On the multilingual safety of large language models. *arXiv preprint arXiv:2310.00905*.
- Ellery Wulczyn, Robert West, Leila Zia, and Jure Leskovec. 2016. Growing wikipedia across languages via recommendation. In *Proceedings of the 25th International Conference on World Wide Web*, pages 975–985.
- Jiebin Zhang, Eugene J Yu, Qinyu Chen, Chenhao Xiong, Dawei Zhu, Han Qian, Mingbo Song, Xiaoguang Li, Qun Liu, and Sujian Li. 2024. Retrieval-based full-length wikipedia generation for emergent events. *arXiv preprint arXiv:2402.18264*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Yang Zhong. 2021. Wikibias: Detecting multi-span subjective biases in language. Master’s thesis, The Ohio State University.

A Prompt for generating neutral Wiki-style sentence

Prompt for generating Wiki-style sentence

For each query message, remove framing bias and epistemological bias and do not add any extra content from your own knowledge.

Framing bias: subjective words or one-sided words, revealing the author's stance in a particular debate.

Epistemological bias: propositions that are either commonly agreed to be true or false and that are subtly presupposed, entailed, asserted or hedged in the text.

Here are some examples:.....

Provide only the Output as:
<pad>output</pad>

B Prompt for selecting the most relevant chunk to the given section content

Prompt for selecting most relevant chunk

For each query text, find out whether the given piece of text is relevant or not.

start(*)

.....

end(*)

Evaluate whether the chunk between start(*) and end(*) is relevant to the given section content. A quality chunk should meet the following criteria: a) It should provide relevant information as compared with the content, b) it should be well-written.

Provide the output in the following format:

–Yes/No

– Confidence score: <score>

C Groups in informativeness: FA category

Given the high standard deviation observed in the informativeness metric for c_{old} , it is worthwhile to explore whether the improvements in content c_{new} compared to c_{old} are uniform across all the articles. We categorize the informativeness scores for c_{old} into three ranges based on their distribution and record the corresponding scores for the same sections in c_{new} . Table 5 displays the informativeness scores for both c_{old} and c_{new} across these three groups. It is clear that the informativeness has

improved in c_{new} compared to c_{old} in each group, mirroring the results shown in Table 5.

D Human assessment for generated text

The average assessment score for each evaluator is tabulated in Table 6. Further, the evaluation is shown for FA and non-FA articles separately.

E Old and newly generated content: sample sections

The Hindi output for the FA article is generated using WikiTransfer, and both the Hindi content and its English translation are displayed in Figure 1 and Figure 2. Similarly, for the C-class article, the new Hindi content is first pooled into text from external resources using the RAG method, followed by the NPOV correction and the WikiTransfer framework. The corresponding Hindi output and its English translation for this sample section are presented in Figure 3 and 4, respectively.

Type	Group1 (0–50)	Group2 (50–100)	Group3 (100 and more)
<i>C_{old}</i>	18.6 (13.78)	71.00 (13.92)	177.18 (71.38)
<i>C_{new}</i>	61.79 (46.55)	108.22 (50.07)	235.18 (138.97)

Table 5: Automatic evaluation: mean and (standard deviation) of the metric informativeness divided into ranges of scores for the articles that belong to FA quality class.

Existing Hindi	Newly generated Hindi
<p>फ्रीडा पिंटो मुंबई में सिल्विया पिंटो, जो सेंट जॉन यूनिवर्सल हाई स्कूल (गुडगाँव) में प्रधानाचार्य हैं और फ्रेडरिक पिंटो, जो की बड़ोदा बैंक के एक वरिष्ठ शाखा प्रबंधक हैं उनके घर पैदा हुई। फ्रीडा पिंटो के पिता नीरुड़े से और उनकी माता डेरेबेल से हैं, यह दोनों कस्बे मंगलोर में हैं। उनका परिवार मंग्लोरीयन कैथोलिक मूल का है, और एक इंटरव्यू में, पिंटो ने कहा है कि वह पूरी तरह से शुद्ध भारतीय हैं, लेकिन उनका परिवार कैथोलिक है। उनकी बड़ी बहन शेरोन पिंटो NDTV समाचार चैनल पर एक सहायक निर्माता है। पिंटो ने मलाड में सेंट जोसेफ स्कूल के कर्मल में अभ्यास किया और अंग्रेजी साहित्य में कला स्नातक (BA) की उपाधि सेंट जेवियर्स कॉलेज, मुंबई से ली। वर्तमान समय में वह मुंबई के एक उपनगर, मलाड के औरलेम इलाके में रहती है। वह भारतीय शास्त्रीय नृत्य के विभिन्न रूपों के साथ साथ साल्सा में भी प्रशिक्षित हैं। वह अपने पूर्व प्रचारक और आठ साल के प्रेमी रोहन अन्ताओ के साथ संलग्न थी, लेकिन 2009 की शुरुआत में स्टार बनने के बाद उन्होंने यह सगाई तोड़ दी। आजकल वह अपने 'स्लमडॉग मिलियनेर' फिल्म के सह अभिनेता, देव पटेल के साथ डेटिंग कर रही है, जो की उनसे उम्र में छह साल छोटे है। वह पीपल People पत्रिका की सबसे खूबसूरत लोगों की सूची में और दुनिया की बेहतरीन सजने वाली महिलाओं की सूची में भी शामिल हो चुकी हैं।</p>	<p>फ्रीडा पिंटो मुंबई में सिल्विया पिंटो, जो सेंट जॉन यूनिवर्सल हाई स्कूल (गुडगाँव) में प्रधानाचार्य हैं और फ्रेडरिक पिंटो, जो की बड़ोदा बैंक के एक वरिष्ठ शाखा प्रबंधक हैं उनके घर पैदा हुई। फ्रीडा पिंटो के पिता नीरुड़े से और उनकी माता डेरेबेल से हैं, यह दोनों कस्बे मंगलोर में हैं। उनका परिवार मंग्लोरीयन कैथोलिक मूल का है, और एक इंटरव्यू में, पिंटो ने कहा है कि वह पूरी तरह से शुद्ध भारतीय हैं, लेकिन उनका परिवार कैथोलिक है। उनकी बड़ी बहन शेरोन पिंटो NDTV समाचार चैनल पर एक सहायक निर्माता है। पिंटो ने मलाड में सेंट जोसेफ स्कूल के कर्मल में अभ्यास किया और अंग्रेजी साहित्य में कला स्नातक (BA) की उपाधि सेंट जेवियर्स कॉलेज, मुंबई से ली। वर्तमान समय में वह मुंबई के एक उपनगर, मलाड के औरलेम इलाके में रहती है। वह भारतीय शास्त्रीय नृत्य के विभिन्न रूपों के साथ साथ साल्सा में भी प्रशिक्षित हैं। वह अपने पूर्व प्रचारक और आठ साल के प्रेमी रोहन अन्ताओ के साथ संलग्न थी, लेकिन 2009 की शुरुआत में स्टार बनने के बाद उन्होंने यह सगाई तोड़ दी। आजकल वह अपने 'स्लमडॉग मिलियनेर' फिल्म के सह अभिनेता, देव पटेल के साथ डेटिंग कर रही है, जो की उनसे उम्र में छह साल छोटे है। वह पीपल People पत्रिका की सबसे खूबसूरत लोगों की सूची में और दुनिया की बेहतरीन सजने वाली महिलाओं की सूची में भी शामिल हो चुकी हैं। के लिए काम करती हैं। पिंटो की परवरिश उत्तरी मुंबई के मलाड उपनगर में हुई थी। वह पहली बार एक अभिनेत्री बनना चाहती थी जब वह पाँच साल की थी, अक्सर अपने बचपन के दौरान टेलीविजन अभिनेताओं के कपड़े पहनती थी और उनकी नकल करती थी। बाद में उन्होंने 1994 की मिस यूनिवर्स प्रतियोगिता में सुष्मिता सेन की जीत से प्रेरित होने को याद करते हुए कहा कि देश को वास्तव में उन पर गर्व था, और मुझे लगा, एक दिन, मैं भी ऐसा ही करना चाहती हूँ। पिंटो ने सेंट के कर्मल में भाग लिया। मलाड, उत्तरी मुंबई में जोसेफ स्कूल और फिर सेंट में अध्ययन किया। जेवियर्स कॉलेज, मुंबई, फोर्ट, दक्षिण मुंबई। उनका प्रमुख अंग्रेजी साहित्य में था, जिसमें मनोविज्ञान और अर्थशास्त्र में नाबालिग था। कॉलेज में, उन्होंने शौकिया रंगमंच में भाग लिया, लेकिन 2005 में स्नातक होने तक अभिनय और मॉडलिंग के कार्यों को अस्वीकार कर दिया। कम उम्र से ही अभिनय में उनकी रुचि के बावजूद, पिंटो ने यह तय नहीं किया था कि कॉलेज में मॉन्स्टर (2003) देखने तक कौन सा करियर लेना है।</p>

Figure 1: An example of existing and WikiTransfer generated new content – a sample section that belongs to FA quality – Hindi version.

Existing Hindi (English translation)	Newly generated Hindi (English translation)
<p>Freida Pinto was born in Mumbai to Sylvia Pinto, who is the principal at St. John's Universal High School (Gurgaon) and Frederick Pinto, a senior branch manager at Baroda Bank. Freida Pinto's father is from Neerude and her mother is from Derebail, both towns in Mangalore. His family is of Mangalorean Catholic origin, and in an interview, Pinto has stated that he is a fully pure Indian, but his family is Catholic. Her elder sister Sharon Pinto is an assistant producer on the NDTV news channel. Pinto studied at St. Joseph's School Carmel in Malad and graduated with a Bachelor of Arts (BA) in English literature from St. Xavier's College, Mumbai. Currently, she lives in the Orlem area of Malad, a suburb of Mumbai. She is trained in various forms of Indian classical dance as well as Salsa. She was engaged to Rohan Antao, her former publicist and boyfriend of eight years, but broke off the engagement after she became a star in early 2009. She is currently dating her Slumdog Millionaire co-star, Dev Patel, who is six years her junior. She has also appeared on People magazine's list of the most beautiful people and on the list of the world's best-dressed women.</p>	<p>Freida Pinto was born in Mumbai to Sylvia Pinto, who is the principal at St. John's Universal High School (Gurgaon) and Frederick Pinto, a senior branch manager at Baroda Bank. Freida Pinto's father is from Neerude and her mother is from Derebail, both towns in Mangalore. His family is of Mangalorean Catholic origin, and in an interview, Pinto has stated that he is a fully pure Indian, but his family is Catholic. Her elder sister Sharon Pinto is an assistant producer on the NDTV news channel. Pinto studied at St. Joseph's School Carmel in Malad and graduated with a Bachelor of Arts (BA) in English literature from St. Xavier's College, Mumbai. Currently, she lives in the Orlem area of Malad, a suburb of Mumbai. She is trained in various forms of Indian classical dance as well as Salsa. She was engaged to Rohan Antao, her former publicist and boyfriend of eight years, but broke off the engagement after she became a star in early 2009. She is currently dating her Slumdog Millionaire co-star, Dev Patel, who is six years her junior. She has also appeared on People magazine's list of the most beautiful people and on the list of the world's best-dressed women. working for. Pinto grew up in the Malad suburb of North Mumbai. She first wanted to be an actress when she was five years old, often dressing up and mimicking television actors during her childhood. She later recalled being inspired by Sushmita Sen's victory at the 1994 Miss Universe pageant, saying that the country was \"really proud of her, and I thought, one day, I want to do the same.\" Pinto attended the Carmel of Sts. Joseph School in Malad, North Mumbai, and then studied at St. Xavier's College, Mumbai, Fort, South Mumbai. Her major was in English literature, with minors in psychology and economics. In college, she participated in amateur theatre but declined acting and modelling assignments until graduating in 2005. Despite his interest in acting from an early age, Pinto did not decide which career to pursue until watching Monster (2003) in college.</p>

Figure 2: An example of existing and WikiTransfer generated new content – a sample section that belongs to FA quality – English version of Figure 1.

Existing Hindi	Newly generated Hindi
<p>1863 में गुरु विरजानंद के पास अध्ययन पूर्ण होने के बाद लगभग बीस वर्षों के कार्यकाल में दयानंद सरस्वती की हत्या व अपमान के लगभग 44 प्रयास हुए। जिसमें से 17 बार विभिन्न माध्यमों से विष देकर प्राण हरण के प्रयास हुए।</p>	<p>1863 में गुरु विरजानंद के पास अध्ययन पूर्ण होने के बाद लगभग बीस वर्षों के कार्यकाल में दयानंद सरस्वती की हत्या व अपमान के लगभग 44 प्रयास हुए। जिसमें से 17 बार विभिन्न माध्यमों से विष देकर प्राण हरण के प्रयास हुए। दयानंद को अपनी जान से मारने के कई असफल प्रयासों का सामना करना पड़ा। उनके समर्थकों के अनुसार, उन्हें कुछ मौकों पर जहर दिया गया था, लेकिन हठ योग के नियमित अभ्यास के कारण वे इस तरह के सभी प्रयासों से बच गए। एक कहानी बताती है कि हमलावरों ने एक बार उसे नदी में डुबोने का प्रयास किया था, लेकिन दयानंद ने हमलावरों को नदी में घसीटा, हालांकि उसने डूबने से पहले उन्हें छोड़ दिया। एक अन्य विवरण में दावा किया गया है कि उन पर मुसलमानों ने हमला किया था जो गंगा पर ध्यान देते समय इस्लाम की आलोचना से नाराज थे। उन्होंने उसे पानी में फेंक दिया लेकिन दावा किया जाता है कि उसने खुद को बचाया क्योंकि उसके प्राणायाम अभ्यास ने उसे हमलावरों के जाने तक पानी के नीचे रहने दिया। मैं सत्य का शिष्य हूँ। मैं हमेशा सच बोलूंगा। अगर ऐसा करने में मैंने किसी को चोट पहुंचाई है तो गलती मेरी नहीं है। मैं सच को कभी नहीं छोड़ूंगा। उन्होंने कहा, वह अपने शब्दों पर खरे उतरे। उनके उपदेशों के कारण, कई लोग दयानंद के अनुयायी बन गए। उनमें से कई लोगों ने उनके शिष्य बनने का फैसला किया। न केवल हिंदुओं के बीच बल्कि मुसलमानों और ईसाइयों के बीच भी उनके समर्थक थे। उनमें रेवरेंड स्कॉट और सैयद अहमद खान उल्लेखनीय हैं। हालांकि, कई लोगों ने उनका विरोध किया। वे उसे अपने हितों के लिए खतरा मानते थे। वे उसे धमकी देते थे; उसके भोजन और दूध को जहर देते थे और अक्सर हिंसा के साथ उसका सामना करते थे। उन्होंने शांति से जवाब दिया और उन लोगों के प्रति दया दिखाई जो उन्हें नुकसान पहुंचाने की कोशिश कर रहे थे।</p>

Figure 3: An example of existing and our framework generated new content – a sample section that belongs to C quality – Hindi version.

Existing Hindi (English translation)	Newly generated Hindi (English translation)
In 1863, after completing his studies with Guru Virajanand, there were about 44 attempts on Dayanand Saraswati's life over a period of about twenty years. Out of which, 17 attempts were made to kill by poisoning through various means.	In 1863, after completing his studies with Guru Virajanand, there were about 44 attempts on Dayanand Saraswati's life over a period of about twenty years. Out of which, 17 attempts were made to kill by poisoning through various means. Dayanand faced several unsuccessful attempts to kill himself. According to his supporters, he was poisoned on a few occasions, but he escaped all such attempts due to his regular practice of Hatha yoga. One story states that the attackers had once tried to drown her in the river, but Dayanand dragged the attackers into the river, though he abandoned them before drowning. Another account claims that he was attacked by Muslims who were outraged by criticism of Islam while meditating on the Ganges. They threw him into the water but it is claimed that he saved himself as his pranayama practice left him underwater until the attackers left. "I am a disciple of truth. I will always tell the truth. If I have hurt someone in doing so, it is not my fault. I will never give up on the truth." He lived up to his words. Due to his teachings, many became followers of Dayanand. Many of them decided to become his disciples. He had supporters not only among Hindus but also among Muslims and Christians. Notable among them are Reverend Scott and Syed Ahmad Khan. However, many people opposed it. They considered it a threat to their interests. They threatened her; poisoned her food and milk and often confronted her with violence. They responded calmly and showed kindness to those who were trying to harm them.

Figure 4: An example of existing and our framework generated new content– a sample section that belongs to C quality – English version of Figure 3.

Type	Informativeness		Readability		Coherence	
	FA	non-FA	FA	non-FA	FA	non-FA
Evaluator 1	2.9	2.77	2.8	2.5	2.5	2.43
Evaluator 2	2.2	2.3	2.1	2.5	1.8	2.17
Evaluator 3	3	2.93	2.4	2.33	2.3	2.0
Evaluator 4	2.6	2.7	2.6	2.7	2.6	2.53
Evaluator 5	2.4	2.37	1.9	2.17	2.0	2.03
Evaluator 6	3	3	2.7	2.67	2.8	2.8
Evaluator 7	2.6	2.67	2.7	2.6	2.4	2.3

Table 6: Human evaluation on the generated machine-translated Hindi content based on three metrics – informativeness, readability, coherence.