

MaCmS: Magahi Code-mixed Dataset for Sentiment Analysis

Priya Rani¹, Gaurav Negi¹, Theodorus Fransen², John P McCrae¹

¹ Data Science Institute, University of Galway, Ireland

² Università Cattolica del Sacro Cuore, Milan, Italy

{priya.rani, gaurav.negi, john.mccrae}@insight-centre.org
theodorus.fransen@unicatt.it

Abstract

The present paper introduces new sentiment data, MaCmS, for Magahi-Hindi-English (MHE) code-mixed languages, where Magahi is a less-resourced minority language. This dataset is the first Magahi-Hindi-English code-mixed dataset for sentiment analysis tasks. Further, we provide a linguistic analysis of the dataset to understand the structure of code-mixing and a statistical study to understand the language preferences of speakers with different sentiment categories. With these analyses, we also train baseline models to evaluate the dataset's quality.

Keywords: Magahi, Code-mixing, Less-resourced language, Sentiment Analysis

1. Introduction

Different people can perceive a word differently, and these differences are rarely distinctive; however, these are often linked with social factors, such as age, gender, race, geography and more inferable characteristics, such as political and cultural attitudes (Yang and Eisenstein, 2017). Sociolinguists and psychologists have been studying these variations in the lexicons and the language from the 50's (Fischer, 1958; Labov, 1963). The advent of social media has created extensive opportunities to explore these variations automatically. In multilingual societies, particularly on social media, code-mixing, which involves using multiple languages within a single conversation or text, is a well-known phenomenon. This practice has expanded due to the global reach of social media, making it a valuable source of texts for less-resourced languages. One of the applications of this data is sentiment analysis, which seeks to understand people's emotions and attitudes towards various subjects. However, this becomes a challenge for low-resourced languages, primarily because text data in these languages often includes a mix of different scripts and languages.

Sentiment analysis is commonly regarded as a task involving categorising text into one of three categories: positive, negative, or neutral, as mentioned in a study by Phani et al. (2016). In recent years, due to the proliferation of social media platforms like YouTube and Twitter, sentiment analysis has expanded into a broader range of applications in various fields. Multinational companies can use it to analyse customer feedback and reviews. In politics, it can gauge public sentiment towards candidates and policies. For social media platforms, it can be used to monitor and filter out offensive content. Sentiment analysis not only reveals the mood and emotional state of the

speaker but also provides insights into cultural and political attitudes. This information is invaluable for organisations, governments, and businesses in understanding public perception and sentiment towards their products, policies, or initiatives. The brevity and informality of social media posts, like tweets and YouTube comments, present unique challenges and opportunities for sentiment analysis. Short text analysis can be more challenging than longer texts, as less context might be available for determining sentiment. Moreover, issues like code-switching and the unavailability of resources for less-resourced languages can complicate the tasks further (Kenyon-Dean et al., 2018).

While sentiment analysis as a field has been expanding, and numerous systems have made remarkable progress in setting new performance standards, the effectiveness of sentiment prediction in the context of code-mixed data still needs to be improved. This limitation is primarily attributed to the variability in language availability and the quality of training data, which directly impacts the precision of sentiment analysis. Sentiment analysis for Indian languages, especially in code-mixed settings, is still relatively nascent (Jose et al., 2020; Chakravarthi et al., 2020a,b; Rani et al., 2020). The significant difference in style of language, orthography (Chakravarthi et al., 2019) and grammar used in tweets presents specific challenges for the Indian languages code-mixed data especially when it comes to low-resource languages like Magahi. Therefore, this work presents a sentiment analysis dataset for Magahi-Hindi-English code-mixed language. Our major contributions to the study are as follows:

- MaCmS: an annotated Magahi (MAG), Hindi (HIN), English (ENG) code-mixed dataset for sentiment analysis. To the best of our knowledge, this is the first Magahi code-mixed dataset for sentiment analysis.

- A linguistic analysis of the structure of code-mixing between two closely related languages, Magahi and Hindi. Which helps understand when and where the code-mixing is happening.
- A statistical analysis of the dataset based on the language preference used by the speakers, which indicates the emotions and attitude of the speakers and the sentiment they show.
- We also provide some baseline models for sentiment analysis at the sentence and language-specific span levels.

2. Literature Review

There is abundant research work on sentiment analysis. These works study the diverse information about individuals on social media platforms. The primary goal of these studies is to gain insights into the emotions experienced by individuals, which play a fundamental role in personal development, intellectual growth, and human understanding (Xia et al., 2020). Even though the past decade has taken the study of sentiment analysis to a more advanced level, it still lacks in the case of less-resourced languages, and the performance of state-of-the-art is yet to be surpassed in different multilingual settings (Ali et al., 2021).

Several resources have been created for sentiment analysis over the period. Yu et al. (2020) introduced a Chinese single- and multi-modal sentiment analysis dataset, CH-SIMS, which contains 2,281 refined video segments in the wild with both multimodal and independent unimodal annotations. XED, a multilingual fine-grained emotion dataset which consists of human-annotated 25 K Finnish and 30 K English sentences, as well as projected annotations for 30 additional languages, providing new resources for many low-resource languages (Öhman et al., 2020). Patwa et al. (2020) organised a shared task on sentiment analysis, leading to the release of the Hindi-English and Spanish-English language pairs dataset.

The formulation of practical guidelines for sentiment analysis by Mohammad (2016) spurred the creation of numerous sentiment analysis corpora. This surge in corpus development was also notable in the Indian subcontinent characterized by rich language diversity the language diversity, the efforts led to the creation of sentiment corpora for many of the less-resourced languages. Ali et al. (2021) created a Sindhi subjective lexicon with the help of existing English resources. This resource consists of opinion words, SentiWordNet, a Sindhi-English bilingual dictionary and a collection of Sindhi Modifiers. Each opinion word is paired with a positive or negative sentiment score.

Chakravarthi et al. (2020b) developed a Tamil-English sentiment analysis corpus with 15,744 comment posts from YouTube. The authors also provide the baseline models for the gold standard dataset. Ehsan et al. (2023) developed BiLSTM-based models to train sentiment analysis classifiers for Tamil and Tulu code-switched datasets by deploying contextualized word embeddings at input layers of the models. The F1-score of the models are 0.2877 and 0.5133 for Tamil and Tulu datasets, respectively.

Patra et al. (2015) organised a shared task competition on sentiment analysis of code-mixed data pairs for Hindi-English and Bengali-English. The best-performing team used SVM for sentence classification. Kulkarni et al. (2021) published L3CubeMahaSent, the first publicly available Marathi sentiment analysis dataset. It consists of 16,000 distinct tweets classified into three classes: positive, negative, and neutral, with extensive baseline models using CNN, LSTM, ULM-FiT, and BERT. Gokani and Mamidi (2023) describe the Gujarati Sentiment Analysis Corpus (GSAC), sourced from Twitter and manually annotated by native speakers. The paper discusses the creation process of the dataset and provides extensive baseline experiments with the highest F1-score of 0.66 from the IndicBert model.

3. MaCmS Dataset

3.1. Data Creation and preprocessing

YouTube is a popular and readily available source for collecting Magahi corpus. We gathered the data from publicly available comments on YouTube as some Magahi speakers are highly active on YouTube. For this task, we selected two YouTube channels, namely, 'Magadhi Boys'¹ and 'Magadh Music'². As concluded in the study conducted by Rani et al. (2022), both channels show significant activity, featuring videos spanning various topics including folklore, mythology, society, politics, environment, and entertainment. This presents us with an opportunity to explore the sentiments and attitudes of Magahi speakers. Moreover, being a minority and a less-resourced language, Magahi has a minimal presence on social media, thus giving us less opportunity to collect digital data. Further, we removed the duplicate comments from the data and labelled the polarity of the comments at the sentence level. Secondly, We tokenised the comments into language-specific spans and tagged these spans with their polarity.

¹<https://www.youtube.com/channel/UCvH5PbwK8I3lyRSQSjsqYwQ>

²<https://www.youtube.com/c/MagadhMusic1>

3.2. Data Annotation

The dataset aimed to get the polarity of the comments for sentiment analysis in closely related code-mixed text for low-resourced settings. At the same time, we are also using the dataset to get language-specific features from the language-specific span and its polarity. This dataset includes annotations from four annotators. Two of the four annotators were female, and the other two were male. Among the two female annotators, one is a trained linguist, and the other is a computer science student. Among the two male annotators, one is a language student, and the other is a mathematics student. However, all four annotators' mother tongue was Magahi, and they were fluent in all three languages involved in the dataset. They are aged 19-28 and belong to middle to upper-middle-class families. The four annotators originally belong to Bihar, India, where Magahi is one of the dominant languages. We asked the annotators to annotate the polarity of the text at two levels: sentence level and language-specific span, discussed below:

- Sentence Level: We label a comment's polarity at the sentence level.
- Span level: Where each comment is tokenised with language-specific span or, in simple terms, we tokenise the code-mixed text into language-specific tokens from its code-mixing points. The tokenisation is done manually with the help of native speakers.

Number	Sentence	Translation	Sentiment
Sent: 1	Bahut sundar bahut wait kara hi apne vedio ke	Very nice I am eagerly waiting for your video	Positive
Sent: 2	तौरा दोनो के सामने सलमान खान शाहरुखखान अमीर खान अजय देवनसब फेल है जी हमको बिहारी होने पे इतना गरब है कि बता नय सक्कीयो हा जी तौर दोनो के वीडियो देखो हीयो तो मन खुश हो जाहो	Salman Khan, Shahrukh Khan and Ajay Devgan are nothing in front of you both; I am so proud to be a Bihari that I can't express. Moreover, I am delighted whenever I see your videos	Positive
Number	Span and (language)	Translation	Sentiment
Sent: 1 Span1	Bahut sundar bahut (Hindi)	Very nice	Positive
Span2	wait (English)	wait	Neutral
Span3	kara hi apne vedio ke (Magahi)	For your video	Neutral
Sent: 2 Span1	तौरा दोनो के सामने सलमान खान शाहरुखखान अमीर खान अजय देवनसब फेल है जी (Magahi)	Salman Khan, Shahrukh Khan and Ajay Devgan are nothing in front of you both	Positive
Span2	हमको बिहारी होने पे इतना गरब है कि (Hindi)	I am so proud to be bihari that	Positive
Span3	बता नय सक्कीयो हा जी तौर दोनो के वीडियो देखो हीयो तो मन खुश हो जाहो (Magahi)	I can't express. Moreover, I am delighted whenever I see your videos.	Positive

Figure 1: Demonstrations of the annotation with examples.

3.3. Annotation Guidelines

Positive Sentiment (POS): A sentence or a language-specific span of a comment is labelled with positive polarity with the tag POS when the text reflects the following criteria:

- The speaker uses positive language to express support, admiration, positive attitude, forgiveness, fostering, success, and a positive emotional state.
- The choice of annotation should remain independent of agreeing or disagreeing with the speaker's opinions. Evaluators should focus on assessing the language rather than judging the views expressed.
- The text contains directly or indirectly implied indications that the speaker is experiencing positive emotions such as happiness, admiration, relaxation, and forgiveness.

NEG: Negative Sentiment

- The speaker uses negative language, for example, expressions of criticism, judgement, negative attitude, questioning validity/competence, failure, and negative emotion.
- There is an explicit or implicit clue in the text suggesting that the speaker is in a negative state.
- The speaker is using expressions of sarcasm, ridicule or mockery.

NEU: Neutral Sentiment

- The speaker is using neither positive nor negative language
- There is no explicit or implicit clue of the speaker's emotional state indicating the speaker feels positively or negatively

MIX: Mixed sentiment

- The speaker is using positive and negative language in part.
- There are explicit and implicit clues in the text suggesting that the speaker has positive and negative feelings.

3.4. Workflow of annotation

The annotation was accomplished in three phases. In the first phase, we gave each annotator a sample of 500 sentences and 100 spans. After the annotation, we manually evaluated the annotation.

We made some subtle changes in the annotation guidelines, including a more detailed explanation of the tag and providing examples as references for each annotation guideline. The initial guidelines were motivated by the practical guide to sentiment analysis discussed by [Mohammad \(2016\)](#). Following the initial annotation phase, we conducted a feedback session with the annotators to enhance the quality of annotations and ensure a thorough understanding of the annotation guidelines. The second phase of the annotation was again conducted with another set of samples with 500 sentences and 100 spans, where we found much improvement in the annotation task. At last, when we were satisfied with the annotators, we gave them the data for the final round of annotation with a set of 1000 sentences and 500 spans one at a time. With each batch of annotation, we manually evaluated the sentiment of each annotator, and in disagreement, we followed the rules mentioned below to get the final decision.

- When there is a disagreement between all the annotators, we re-annotated the text by ourselves.
- If there is disagreement by one annotator but the other three annotators agree on a tag, then the final tag of the text would be the tag selected by the majority. For example, if three of the annotators tagged a text or a span with POS and one of the annotators tagged it as NEU, then the final tag would be POS.
- If two annotators agree and the other two have two different opinions about the polarity of the text, then the final tag was given with majority voting. For example, if a text is labelled with POS by two annotators while the third annotator labelled it as NEG and the fourth annotator marked it as NEU, then in this case, the final label would be POS.
- When the agreement between the annotators is distributed equally, for example, if two annotators labelled the text with POS and the other two annotators labelled it as NEU, we sent the text for re-annotation.

4. Data Analysis

4.1. Inter-annotator Agreement

After each annotation phase, we calculated the inter-annotator agreement to check the annotation process's reliability and ensure the dataset quality. We used Krippendorff's α to estimate the agreement between the annotators. The agreement score helped us visualize the improvement we need to make in the annotation guidelines to

improve the quality of the annotated dataset. Table 1 shows the obtained IAA between the annotators at each phase. Considering the nature of the task is very subjective, we achieved a fair final agreement score between the annotators, that is, 0.78 for sentence-level annotation and 0.76 for span-level annotation.

Phase	Sentence-level	Span-level
Phase-1	0.67	0.69
Phase-2	0.72	0.76
Final	0.78	0.76

Table 1: Krippendorff's α for inter-annotator agreement of the MHE sentiment analysis dataset

4.2. Data Statistics

We collected 11,000 comments from the YouTube channels mentioned. However, we discarded the comments in English or Hindi, and at last, we were left with 5663 comments. Out of these, 5,000 comments were labelled for sentence sentiment analysis. We tokenised the comments for language-specific span. Due to the annotators' time constraints, we could annotate only 750 sentences for span-level sentiment analysis, which resulted in 2642 spans. Table 2 gives brief statistics of the dataset.

MHE	Number
Sentences	5000
Span sentences	750
Total span	2642

Table 2: Statistics of our YouTube Magahi-Hindi-English (MHE) code-mixed dataset for sentiment analysis

The distribution of the sentiment tags for sentence-level analysis is demonstrated in Figure 2. We find that the distribution of positive and negative tags is mostly balanced, with 37.4% and 33.8%, respectively. However, the neutral and mixed sentiment percentage is low, with 16.2% and 12.6%, respectively. In contrast, the rate of neutral tags in the language-specific span exceeds and consists of 50.2% of the whole dataset, as shown in Figure 3. We also find that the span with positive tags acquires the second position with 35.0% of the entire sentiment dataset, followed by negative tags with 9.8% and mixed tags with 5%.

5. Baselines Experiments

This section will briefly describe the various models used for the baseline experiments. We conduct




Figure 2: Distribution of the tags in the MHE dataset for sentence-level sentiment analysis.




Figure 3: Distribution of the tags in the MHE dataset for language-specific span-level sentiment analysis.

a very limited pre-processing that removes any URLs from the data before or after dividing the data into the training, validation, and test sets with the distribution of 70%, 20% and 10%, respectively. However, no further pre-processing steps, such as casing or data normalisation, are conducted to make the baseline experiments more robust.

5.1. mBERT

We have used the multilingual pre-trained BERT (Devlin et al., 2019; Turc et al., 2019) model to fine-tune our first baseline model. It has an embedding dimension of 768. We have trained our model for 10 epochs on the training data set and have used a stepped LR scheduler for the learning rate schedule. The learning rate is set to $2e-5$. Based on the Hugging Face implementation, we have used the below equation 1 as a warm-up step definition for training the model. Here, ‘ r ’ is the tuneable parameter, which defines the percentage of data used to define the step size while training.

$$W_{steps} = \frac{(len(training_set) \times epochs_{training})}{batchsize_{training} \times r} \quad (1)$$

5.2. XLM-R

We use XLM-RoBERTa (XLM-R) (Conneau et al., 2020) to fine-tune the second model. As XLM-R is pre-trained on 100 languages, which includes Indian languages such as Hindi and Hindi being one of the closely related languages to Magahi, it would provide some aid to train the model on the Magahi-Hindi-English dataset for sentiment analysis. As shown in the previous research (Winata et al., 2021) that the XLM-R model performs better than other multilingual models, we investigate the effectiveness of XLM-R in less-resource language settings. We used pre-trained models from Hugging Face³. We put a fully connected classifier on each model and trained the model with a learning rate of $1e-5$ with a decay of 0.1 and batch size of 32 for 10 epochs.

5.3. GenMA

The Generative Morphemes with Attention (GenMA) model (Goswami et al., 2020) is a sentiment analysis model trained to classify sentiment based on newly generated artificial character sequences termed as artificial morphemes. This model combines two convolution layers with one max pooling layer, a BiLSTM layer and an attention layer. As the model generates new character sequences irrespective of language and orthographic features, it helps to capture both language-specific and code-mixed language features. Thus, this model is used to train sentiment classifiers on less-resource languages.

We have used 32 filters, each with a kernel size of 3. The max-pooling size is 3. The hidden size of LSTM units is kept to 100. The dense layer has 32 neurons and a 50 percent dropout. The Adam optimizer (Kingma and Ba, 2015) trains the model with the default learning set to 0.0001. The batch size is set to 10. We have used the Relu activation function (Nair and Hinton, 2010) and tanh activation function (Kalman and Kwasny, 1992) for the convolution and dense layer, respectively. Categorical cross-entropy loss is used for the multi-class classification.

6. Results

Overall, we see varying performance across the models, with some performing much better than others. We computed the F1 score, Precision and Recall to evaluate the performance of the models on our dataset to account for the imbalance in label distribution (see Figure 3). The evaluation scores of the models are summarised in Table 3 for sentence-level sentiment analysis and Table 4

³<https://huggingface.co/>

for language-specific span-level sentiment analysis.

Models	Precision	Recall	F1 Score
mBERT	0.65	0.71	0.68
XLM-R	0.76	0.75	0.75
GenMA	0.68	0.69	0.68

Table 3: The evaluation results of the models at sentence level sentiment analysis. The bold results reflect the best performance of the model.

Models	Precision	Recall	F1 Score
mBERT	0.55	0.52	0.53
XLM-R	0.50	0.48	0.49
GenMA	0.51	0.50	0.51

Table 4: The evaluation results of the models at span level sentiment analysis. The bold results reflect the best performance of the model.

Looking at the evaluation table, we can observe that XLM-R performed the best for sentence-level sentiment analysis with a F1-score of 0.75, while the mBERT F1-score was the highest for span-level sentiment analysis. The poor performance of the span-level models is due to the lack of training data. These results also evaluate the quality of the dataset. As we know, these large language models are data-hungry. They performed comparatively well in code-mixed and less-resourced settings, concluding that the data annotation quality is good.

7. Discussion

7.1. Statistical Analysis

In order to do a statistical analysis to identify the speakers' language preference, especially in closely related code-mixed datasets, we count the number of each sentiment tag in each language span; see Figure 4. The frequency of the Magahi span is high in the positive and neutral domains. Hindi is more prevalent in negative sentiment, whereas English is more frequent in mixed sentiment. While looking at the distribution of the tags across the language, we can deduce the following points:

1. Code-mixing affects the overall sentiment of the comments (sentence-level sentiment) as we can see in Figure 2 that the percentage of positive and negative sentiment is higher than neutral sentiment in sentence-level sentiment. In contrast, the distribution of neutral sentiment is high in language-specific span.

2. The speakers use more Magahi to express positive or neutral sentiments, which implies that the speakers are trying to leave a positive impression of the matrix language, which is Magahi in this case. It could also depend on the theme of the videos in which the comment is posted. For example, a video on Magahi culture, folklore and movies will have a more positive labelled Magahi span to give listeners a positive impression of the Magahi language and culture.
3. The speakers use Hindi to express negative sentiments. This could be because the speakers emphasise their disagreement towards the non-Magahi community, especially in the videos, which discuss the domination of Hindi as a language and society over the Magahi community.




Figure 4: Distribution of the tags in the MHE dataset for language-specific span-level sentiment analysis.

7.2. Linguistic Analysis

The following section provides a linguistic analysis of the proposed dataset. We conducted the analysis manually to study the linguistic characteristics and understand the fundamental structure of code-mixing in closely related languages.

1. Like most closely related languages, Hindi and Magahi share a lexicon, possibly due to various reasons like phonetic similarity in the spelling across languages. For example, the letter 'u' represents *that* in Magahi and *You* in English. Similarly, many tokens are responsible for creating inconsistent tags. However, when the speakers try to express their strong emotions or gain attention, they code-mix a substantial number of functional words, including Wh-words, adjectives, pronouns, determiners, etc.
- (1) E [language (lang) -"MAG"] Match me aapka dubbing matching ho gya Sirji [lang-"HIN"]

Gloss: this match in your-3SG.HON Dubbing match finish-PST AUX Sirji
Translation: Your dubbing got matched in this match Sirji

In Example 1, the determiner \mathbb{E} is inserted in the utterance. This is when the speaker tries to emphasize a certain object, and in this example, that object is *match*.

- (2) बहुते [lang-“MAG”] सुंदर [lang-“HIN”] अपने मगही मे गीत गइली मन गदगद होगेल।
Gloss: very beautiful you-2SG.HON Magahi in song sing-PST.2SG.HON mind happy be-PST
Translation: You sang a very beautiful song in Magahi. It was mind-blowing.

The Example 2 shows the code-mixing between Magahi and Hindi where a Hindi adjective सुंदर (beautiful) is inserted in a Magahi sentence.

2. Insertion of the marker $-wa$ and the numeral classifier go with the noun and act like a classifier language rather than a noun class like Hindi.
- (3) Ek [lang-“HIN/MAG”] du [lang-“MAG”] go [lang-“MAG”] aur [lang-“HIN”] likhiye [lang-“HIN”]
Gloss: one two go-CLF more write-PRS.2SG.HON
Translation: Write one or two more.

In Example 3 with the insertion of classifier $-go$, the text does not exhibit the number agreement in the noun morphology.

- (4) WOW haste haste [lang-“HIN”] mera [lang-“HIN”] pet-**wa** [lang-“MAG”]fat jayega [lang-“HIN”]
Gloss: WOW laugh-PRS.PROG.1SG my stomach-POSS hurt-PRS.PROG.1SG
Translation: My stomach hurts from laughing

In example 4, the noun *pet* is affixed with the familiarity marker $-wa$, presupposing that the noun is familiar to both the speaker and the listener.

3. When the speakers try to quote somebody or some famous expression, they code-mix as shown in Example 5. In the given example, the speaker recites the quote in Magahi, using the post-position k , thus making the quote in Magahi.

- (5) Once a legend said [lang-“ENG”] “Laura k sarkaar”[lang-“MAG”]
Translation: Once a legend said “evil government”

4. While talking about culture or regional traditions, the speakers often code-mix a lot as they are comfortable talking about their culture and tradition in their mother tongue.

- (6) Murna [label-“NAME”] ke [lang-“MAG”] khasi [label-“NAME”] khaiye [lang-“MAG”] ke [lang-“MAG”] hai [lang-“MAG”] Murna [label-“NAME”] kab [lang-“HIN”] hai [lang-“HIN”]☺
Gloss: Murna GEN mutton eat-PRS.1SG bePRS when bePRS
Translation: We have to get Murna’s party, When is the Murna?

In Example 6, the speaker talks about a ritual called *Murna*. Traditionally, the person doing the rituals is supposed to throw a party, which is a very cultural aspect of Magahi society; therefore, the speaker code-mixes while discussing these cultural aspects. Moreover, this comment is also a perfect example of intra-sentential code-mixing.

5. When the speakers try to express their strong emotion of surprise or any other emotion, they are prone to code-mix between Magahi and English as they insert interjections or exclamation in the text. For example, the insertion of *WOW* at the beginning of the comment in example 4 expresses that the speaker is pleased and has a positive emotion.

- (7) Bhaiya [“HIN/MAG”] hum [“HIN”] kehrhe [“HIN”] hain [“HIN”] **ki** [“HIN/MAG”] E [“MAG”] sarkar [“NAME”] ke [“MAG”] sabhe [“MAG”] kala [“HIN/MAG”] cita [“HIN/MAG”] khol [“MAG”] da [“MAG”].
Gloss: Brother I say-1SG.PROG that-COMP this government evil-doing bring out
Translation: Brother, I am telling you to bring out all the evil-doing of this government.

Similarly, when the complementiser is used to express their strong emotion, the speakers tend to code-switch, as shown in example 7.

From this analysis, we tried pinpointing some of the linguistic properties of code-mixing in the dataset. However, these user-generated texts exhibit other characteristics like borrowing. The scenario where borrowing is taking place is very different from code-mixing. For example, the speakers borrow words from English in the expression

where they find it challenging to find the equivalent words in Hindi. For example:

- (8) ई dislike कोन करलो हे साच हई हमनी अपन भासा काहे भुली
Transliteration: E dislike kaun karlo hae saach hai humni apen bhasa kahe bhuli
This dislike who do-PST AUX true we ours language why forget
Translation: Who disliked it? This is true; why will I forget my language?

In Example 8, the word `dislike` is being borrowed. It does not make any difference in the linguistic structure or the discourse if the word is replaced with an equivalent word. However, in YouTube text, words such as `like`, `subscribe`, `comment`, and `dislike` are easily borrowed from English, especially for Indian language speakers, as some of these words are complicated to type in the Devanagari script due to the complex typing format.

8. Conclusion

In this paper, we described the development process of the Magahi-Hindi-English code-mixed dataset annotated with sentiment. We analysed the speaker's language preference for a specific context with the help of a newly developed dataset. Moreover, we discussed some of the linguistic features that gave us insight into the linguistic structure of the code-mixed text. The analysis concluded that Magahi is used to express positive sentiments more than negative sentiments. The findings contradict earlier research (Agarwal et al., 2017; Rudra et al., 2016; Doğruöz et al., 2021). This discrepancy may be attributed to the phenomenon of code-mixing between closely related languages, namely Magahi and Hindi. Additionally, the dominance of Hindi in regions where Magahi is spoken could influence the choice of language for expressing various emotions.

We have also experimented with the baseline models using deep learning and SOTA transformer models. The results of the baseline models show that the multilingual model is relatively good in less-resourced and code-mixed settings. However, our future work will include model engineering to enhance the performance of SOTA so that we can capture the code-mixed representation well, which helps to capture the speakers' emotions efficiently.

9. Limitations

Due to the limited contents in Magahi, we could only collect a few texts. Some of the content is not publically accessible due to privacy reasons in

the provided YouTube channels, which limits access to the data. To study language preference in closely related code-mixed scenarios, we need more studies to establish a final statement other than quantitative analysis.

10. Ethics Statement

All the data collected has been obtained legally and validly, adhering to the provided guidelines. We developed the dataset under the ethical and legal framework of our university and have followed legal requirements as they apply to fair research use of such data. We have anonymised all comments and will remove any comments from the dataset on request should anything be found non-compliant. We followed proper data statement guidelines (Bender and Friedman, 2018) to annotate the dataset to prevent ethical issues. The data might contain strong language, which might be unsuitable for some applications. We will use the data only for research to preserve the users' privacy.

11. Acknowledgement

This work was supported by Science Foundation Ireland (SFI) as part of Grant Number SFI/12/RC/2289_P2 Insight_2, Insight SFI Centre for Data Analytics and co-funded by Science Foundation Ireland (SFI) as part of grant SFI/18/CRT/6223 (CRT-Centre for Research Training in Artificial Intelligence).

References

- Prabhat Agarwal, Ashish Sharma, Jeenu Grover, Mayank Sikka, Koustav Rudra, and Monojit Choudhury. 2017. *I may talk in english but gaali toh hindi mein hi denge : A study of english-hindi code-switching and swearing pattern on social networks*. In *2017 9th International Conference on Communication Systems and Networks (COMSNETS)*, pages 554–557.
- Wazir Ali, Naveed Ali, Yong Dai, Jay Kumar, Saifullah Tumrani, and Zenglin Xu. 2021. *Creating and evaluating resources for sentiment analysis in the low-resource language: Sindhi*. In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 188–194, Online. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. *Data statements for natural language processing: Toward mitigating system bias and enabling better science*. *Transactions of the Association for Computational Linguistics*, 6:587–604.

- Bharathi Raja Chakravarthi, Mihael Arcan, and John P. McCrae. 2019. [Comparison of Different Orthographies for Machine Translation of Under-Resourced Dravidian Languages](#). In *2nd Conference on Language, Data and Knowledge (LDK 2019)*, volume 70 of *OpenAccess Series in Informatics (OASlcs)*, pages 6:1–6:14, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Bharathi Raja Chakravarthi, Navya Jose, Shardul Suryawanshi, Elizabeth Sherly, and John Philip McCrae. 2020a. [A sentiment analysis dataset for code-mixed Malayalam-English](#). In *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL)*, pages 177–184, Marseille, France. European Language Resources association.
- Bharathi Raja Chakravarthi, Vigneshwaran Muralidaran, Ruba Priyadharshini, and John Philip McCrae. 2020b. [Corpus creation for sentiment analysis in code-mixed Tamil-English text](#). In *Proceedings of the 1st Joint Workshop on Spoken Language Technologies for Under-resourced languages (SLTU) and Collaboration and Computing for Under-Resourced Languages (CCURL)*, pages 202–210, Marseille, France. European Language Resources association.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- A. Seza Doğruöz, Sunayana Sitaram, Barbara E. Bullock, and Almeida Jacqueline Toribio. 2021. [A survey of code-switching: Linguistic and social perspectives for language technologies](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1654–1666, Online. Association for Computational Linguistics.
- Toqeer Ehsan, Amina Tehseen, Sarveswaran Kengatharaiyer, and Amjad Ali. 2023. [Sentiment analysis of code-mixed tamil and tulu by training contextualized elmo representations](#).
- John L Fischer. 1958. Social influences on the choice of a linguistic variant. *Word*, 14(1):47–56.
- Monil Gokani and Radhika Mamidi. 2023. [GSAC: A Gujarati sentiment analysis corpus from Twitter](#). In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 129–137, Toronto, Canada. Association for Computational Linguistics.
- Koustava Goswami, Priya Rani, Bharathi Raja Chakravarthi, Theodorus Fransen, and John P. McCrae. 2020. [ULD@NUIG at SemEval-2020 task 9: Generative morphemes with an attention model for sentiment analysis in code-mixed text](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 968–974, Barcelona (online). International Committee for Computational Linguistics.
- Navya Jose, Bharathi Raja Chakravarthi, Shardul Suryawanshi, Elizabeth Sherly, and John P. McCrae. 2020. A survey of current datasets for code-switching research. In *2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS)*, pages 136–141.
- Barry L Kalman and Stan C Kwasny. 1992. Why tanh: choosing a sigmoidal function. In *[Proceedings 1992] IJCNN International Joint Conference on Neural Networks*, volume 4, pages 578–581. IEEE.
- Kian Kenyon-Dean, Eisha Ahmed, Scott Fujimoto, Jeremy Georges-Filteau, Christopher Glasz, Barleen Kaur, Auguste Lalande, Shruti Bhandari, Robert Belfer, Nirmal Kanagasabai, Roman Sarrazingendron, Rohit Verma, and Derek Ruths. 2018. [Sentiment analysis: It’s complicated!](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1886–1895, New Orleans, Louisiana. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

- Atharva Kulkarni, Meet Mandhane, Manali Likhitkar, Gayatri Kshirsagar, and Raviraj Joshi. 2021. [L3CubeMahaSent: A Marathi tweet-based sentiment analysis dataset](#). In *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 213–220, Online. Association for Computational Linguistics.
- William Labov. 1963. The social motivation of a sound change. *Word*, 19(3):273–309.
- Saif Mohammad. 2016. [A practical guide to sentiment annotation: Challenges and solutions](#). In *Proceedings of the 7th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 174–179, San Diego, California. Association for Computational Linguistics.
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814.
- Emily Öhman, Marc Pàmies, Kaisla Kajava, and Jörg Tiedemann. 2020. [XED: A multilingual dataset for sentiment analysis and emotion detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6542–6552, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Braja Patra, Dipankar Das, Amitava Das, and R.Rajendra Prasath. 2015. [Shared Task on Sentiment Analysis in Indian Languages \(SAIL\) Tweets - An Overview](#).
- Parth Patwa, Gustavo Aguilar, Sudipta Kar, Suraj Pandey, Srinivas PYKL, Björn Gambäck, Tanmoy Chakraborty, Tamar Solorio, and Amitava Das. 2020. [SemEval-2020 task 9: Overview of sentiment analysis of code-mixed tweets](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 774–790, Barcelona (online). International Committee for Computational Linguistics.
- Shanta Phani, Shibamouli Lahiri, and Arindam Biswas. 2016. [Sentiment analysis of tweets in three Indian languages](#). In *Proceedings of the 6th Workshop on South and South-east Asian Natural Language Processing (WSANLP2016)*, pages 93–102, Osaka, Japan. The COLING 2016 Organizing Committee.
- Priya Rani, John P. McCrae, and Theodorus Franssen. 2022. [MHE: Code-mixed corpora for similar language identification](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3425–3433, Marseille, France. European Language Resources Association.
- Priya Rani, Shardul Suryawanshi, Koustava Goswami, Bharathi Raja Chakravarthi, Theodorus Franssen, and John Philip McCrae. 2020. [A comparative study of different state-of-the-art hate speech detection methods in Hindi-English code-mixed data](#). In *Proceedings of the Second Workshop on Trolling, Aggression and Cyberbullying*, pages 42–48, Marseille, France. European Language Resources Association (ELRA).
- Koustav Rudra, Shruti Rijhwani, Rafiya Begum, Kalika Bali, Monojit Choudhury, and Niloy Ganguly. 2016. [Understanding language preference for expression of opinion and sentiment: What do Hindi-English speakers do on Twitter?](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1131–1141, Austin, Texas. Association for Computational Linguistics.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv e-prints*, pages arXiv-1908.
- Genta Indra Winata, Samuel Cahyawijaya, Zihan Liu, Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2021. [Are multilingual models effective in code-switching?](#) In *Proceedings of the Fifth Workshop on Computational Approaches to Linguistic Code-Switching*, pages 142–153, Online. Association for Computational Linguistics.
- Hongbin Xia, Chenhui Ding, and Yuan Liu. 2020. [Sentiment analysis model based on self-attention and character-level embedding](#). *IEEE Access*, 8:184614–184620.
- Yi Yang and Jacob Eisenstein. 2017. [Overcoming language variation in sentiment analysis with social attention](#). *Transactions of the Association for Computational Linguistics*, 5:295–307.
- Wenmeng Yu, Hua Xu, Fanyang Meng, Yilin Zhu, Yixiao Ma, Jiele Wu, Jiyun Zou, and Kaicheng Yang. 2020. [CH-SIMS: A Chinese multimodal sentiment analysis dataset with fine-grained annotation of modality](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3718–3727, Online. Association for Computational Linguistics.

