Sentiment Analysis for Konkani using Zero-Shot Marathi Trained Neural Network Model

Rohit M. Ghosarwadkar[1], Seamus Fred Rodrigues[1], Pradnya Bhagat[1], Alvyn Abranches[1], Pratik D. Korkankar[1,2], Jyoti D. Pawar[1]
[1]Goa Business School, Goa University, India.
[2]Dnyanprassarak Mandal's College and Research Centre, Assagao, Goa, India.
rohitgee23@gmail.com, sea11rod2@gmail.com, dcst.pradanya@unigoa.ac.in, dcst.alvyn@unigoa.ac.in, dcst.pratik@unigoa.ac.in, jdp@unigoa.ac.in

Abstract

Sentiment Analysis plays a crucial role in understanding user opinions in various languages. The paper presents an experiment with a sentiment analysis model fine-tuned on Marathi sentences to classify sentiments into positive, negative, and neutral categories. The fine-tuned model shows high accuracy when tested on Konkani sentences, despite not being explicitly trained on Konkani data; since Marathi is a language very close to Konkani. This outcome highlights the effectiveness of Zero-shot learning, where the model generalizes well across linguistically similar languages. Evaluation metrics such as accuracy, balanced accuracy, negative accuracy, neutral accuracy, positive accuracy and confusion matrix scores were used to assess the performance, with Konkani sentences demonstrating superior results. These findings indicate that zero-shot sentiment analysis can be a powerful tool for sentiment classification in resource poor languages like Konkani, where labeled data is limited. The method can be used to generate datasets for resource-poor languages. Furthermore, this suggests that leveraging linguistically similar languages can help generate datasets for low-resource languages, enhancing sentiment analysis capabilities where labeled data is scarce. By utilizing related languages, zero-shot models can achieve meaningful performance without the need for extensive labeled data for the target language.

Keywords: Sentiment analysis, Marathi, Konkani, Zero-shot learning, Fine-tuning, Low-resource languages.

1 Introduction

Sentiment Analysis (Medhat et al., 2014), the computational task of identifying and categorizing opinions expressed in text, has gained significant traction in recent years due to the proliferation of online content, particularly product reviews and social media interactions. Understanding sentiments is vital for businesses and public services aiming to tailor their products and services to meet customer expectations and enhance user experience. While various sentiment analysis models have been developed for languages with abundant resources, many low-resource languages, including Konkani, remain underserved in this area.

Konkani (Mascarenhas-Keyes, 2020; Sardesai, 2006) is a language primarily spoken along the western coast of India and has its own unique linguistic features. Despite its rich cultural heritage, the lack of digital resources and language technology tools has hindered the development of effective natural language processing (NLP) applications for Konkani. On the other hand, Marathi (Bloch, 1970), a language closely related to Konkani, has seen advancements in sentiment analysis due to better resource availability and research focus.

In this study, we aim to investigate the applicability of a pre-existing Marathi sentiment analysis models to Konkani text. By examining the performance of this model on both Marathi and Konkani datasets, we seek to highlight the discrepancies in accuracy and effectiveness, ultimately demonstrating the necessity for a dedicated Konkani sentiment analysis datasest and models. Our research underscores the importance of developing language-specific tools that can cater to the unique characteristics of low-resource languages, thereby improving sentiment analysis capabilities in diverse linguistic contexts.

The remainder of this paper is structured as follows: Section 2 reviews related work on sentiment analysis in low-resource languages and cross-lingual transfer learning. Section 3 describes the methodology, including the dataset used, the model architecture, and the training process. Section 4 presents the experimental details, comparing the performance of the model before and after fine-tuning. Section 5 discusses the implications of the results and the potential for further improvements in accuracy. Finally, Section 6 concludes the paper, highlighting that the creation of a Konkani-specific labeled dataset and further fine-tuning could significantly increase the model's performance in sentiment classification.

2 Literature Survey

Zero-Shot Learning (ZSL) is a paradigm in natural language processing (NLP) that allows models to perform tasks without prior exposure to specific examples or labeled data. This capability is particularly valuable in sentiment analysis, where traditional models often require extensive training on labeled datasets. ZSL leverages a model's innate linguistic knowledge to generalize across tasks it has not encountered during training, making it adaptable to new languages and contexts (Joshi et al., 2020). Recent advancements in pre-training models like BERT and its multilingual variants have enhanced the zero-shot performance by capturing cross-linguistic knowledge, thus allowing the models to generalize well to unseen languages and domains (Pires, 2019; Conneau, 2019).

Zero-shot classification models can categorize text into various sentiments (positive, negative, neutral) based on their understanding of language relationships rather than relying on task-specific training data. This approach accommodates a broader range of emotional expressions. For instance, models can be prompted to classify sentiments related to mental health by providing specific emotional labels such as joy, sadness, or anxiety. A similar approach is applied in cross-lingual settings, where pre-trained models like Multilingual BERT can transfer knowledge across languages, achieving high performance in tasks like sentiment analysis without task-specific annotations in the target language (Pires, 2019; Devlin, 2018).

Recent studies have demonstrated that zero-shot sentiment analysis can effectively extend to low-resource languages. By pre-training models using multilingual lexicons or using libraries such as IndicNLP (Kunchukuttan, 2020), researchers have achieved superior performance in sentiment classification across languages like Marathi and Konkani. This method reduces the need for extensive sentence-level annotations and allows for effective cross-linguistic transfer (Conneau, 2019). Such models, when pre-trained on diverse language corpora, can generalize across various languages, making them particularly useful for low-resource languages that lack large annotated datasets (Joshi et al., 2020).

Despite its advantages, zero-shot learning faces several challenges. Models may struggle with tasks requiring specialized knowledge or intricate understanding of context, such as sarcasm or culturally embedded sentiments. Additionally, zero-shot models might exhibit limitations in grasping the full context of sentiment-laden statements, particularly when subtlety is involved (Joshi et al., 2020). The effectiveness of zero-shot learning heavily relies on the quality and breadth of the pre-training data, as models trained on general language patterns may not perform well in highly specialized domains.

One innovative approach involves using multilingual sentiment lexicons for pre-training models, which has shown promising results in improving zero-shot performance for low-resource languages (Kunchukuttan, 2020; Conneau, 2019). The performance of zero-shot sentiment analysis models is typically assessed using various metrics such as accuracy, balanced accuracy, and confusion matrix scores, which help in evaluating their generalizability across different languages and domains.

The exploration of zero-shot sentiment analysis in low-resource languages like Marathi and Konkani highlights its potential as a transformative tool in NLP. By leveraging general linguistic knowledge and multilingual resources, researchers can effectively analyze sentiments without extensive labeled data. However, ongoing research is necessary to address the challenges associated with domain specificity and contextual understanding (Joshi et al., 2020).

3 Methodology

We use two pre-trained models on Marathi as base models to perform the sentiment analysis on Konkani and Marathi datasets.

1. deepampatel/roberta-mlm-marathi (Patel, 2023)
2. l3cube-pune/MarathiSentiment (L3Cube-Pune, 2024)

3.1 Dataset and Data Preparation

The dataset consisted of 6,000 random sentences of product reviews from Amazon.com (Hou et al., 2024). The sentences are translated from English to Marathi and Konkani using the IndicTrans translation model (Gala et al., 2023). Each sentence was annotated with one of three sentiment labels: positive, negative, or neutral. The ground truth for these annotations was derived from three independent human annotators, with their individual assessments merged into a final consensus label for each sentence. This process ensured that the sentiment labels accurately reflected the true sentiment expressed in each sentence, providing a reliable baseline for evaluating the performance of the models during both the fine-tuning process (Howard and Ruder, 2018) and subsequent testing. The curated dataset was divided into two parts. Training set: 80 percent of the dataset (4,800 sentences) was used to fine-tune the models. Testing set: 20 percent of the dataset (1,200 sentences) was reserved for evaluating the models' performance. Prior to giving the data as input into the models, the text was tokenized using the default tokenizer for each model.

3.2 Sentiment Prediction without Fine-Tuning

Before fine-tuning the models, we evaluated their initial performance by giving Marathi sentences as input into each pre-trained model. The models were tasked with predicting sentiment (positive, negative, or neutral) for these sentences based solely on their pre-trained knowledge. This served as a baseline for measuring the improvement after fine-tuning. The output sentiments were recorded and the accuracies of the models in the test set were calculated.

3.3 Fine-Tuning the Models

We fine-tuned both models on the Marathi sentiment-labeled dataset to adapt their parameters to the specific task of Marathi sentiment classification.

Fine-tuning was conducted on the Marathi training dataset curated by us, allowing the models to learn and adjust to the sentiment patterns present in Marathi sentences. Once the models were fine-tuned, they were evaluated again on the test set to assess improvements in sentiment classification accuracy. Since the sentences in our dataset were from product reviews, it helped the model understand the intricacies in sentiments expressed in these reviews.

3.4 Zero-Shot Sentiment Prediction on Konkani Sentences

After fine-tuning on Marathi sentences, both models were applied to predict the sentiment of Konkani sentences. This step was a zero-shot learning scenario because the models had not been explicitly trained on Konkani data. The goal was to evaluate the models' ability to generalize the knowledge learned from Marathi and apply it to Konkani.

The Konkani sentences were pre-processed similarly to the Marathi sentences (tokenization and padding), and the predictions were recorded. Since Konkani and Marathi share lexical and syntactic similarities, the models were expected to transfer their understanding of Marathi sentiment to Konkani.

3.5 Evaluation Metrics

For all stages of the experiment, accuracy (Muhammad et al., 2017) and confusion matrix (Heydarian et al., 2022) were used as the primary metrics to evaluate the models' performance. Accuracy was calculated as the percentage of correctly predicted sentiments (positive, negative, neutral) out of the total number of test sentences. The performance of the models was measured at three points: Before fine-tuning (baseline performance on Marathi sentences), After fine-tuning (performance on Marathi sentences after fine-tuning), Zero-shot prediction (performance on Konkani sentences using the fine-tuned models).

4 Experimentation Details

4.1 Initial Sentiment Prediction without Fine-Tuning

The models were applied directly to the test set of 1,200 Marathi sentences, which had been annotated with human sentiment labels (positive, negative, neutral). The predictions made by the models were compared against the human-annotated ground truth, and the accuracy was calculated.

deepampatel/roberta-mlm-marathi achieved an accuracy of 42.3 percent for Marathi test set without fine-tuning. l3cube-pune/MarathiSentiment yielded an accuracy of 24.5 percent for Marathi test set without fine-tuning.

4.2 Sentiment Prediction after Fine-Tuning on Marathi Sentences

After the initial evaluation, both models were fine-tuned on the training set of 4,800 human-annotated Marathi sentences. Following fine-tuning, the models were again tested on the Marathi test set (1,200 sentences) to assess the improvements in sentiment prediction accuracy.

deepampatel/roberta-mlm-marathi achieved an accuracy of 42.75 percent for Marathi test set with fine-tuning. l3cube-pune/MarathiSentiment yielded an accuracy of 48.16 percent for Marathi test set with fine-tuning.

4.3 Performance of Zero-Shot Sentiment Prediction on Konkani Sentences

After fine-tuning both models on Marathi data, we conducted a zero-shot learning experiment by using the fine-tuned models to predict sentiment in Konkani sentences.

deepampatel/roberta-mlm-marathi achieved an accuracy of 44.41 percent on Konkani test set that was fine-tuned on Marathi train set. l3cube-pune/MarathiSentiment yielded an accuracy of 49.16 percent on Konkani test set that was fine-tuned on Marathi train set.

5 Results and Discussions

This study demonstrates the significant impact of fine-tuning on sentiment classification performance for both Marathi and Konkani. After fine-tuning, both models—deepampatel/roberta-mlm-marathi and l3cube-pune/MarathiSentiment—showed improved accuracy in classifying Marathi sentiments, particularly for deepampatel/roberta-mlm-marathi. Additionally, both models displayed even better performance when applied to Konkani sentences in a zero-shot learning setting, with high accuracy in detecting positive sentiments. These results highlight the effectiveness of cross-lingual transfer learning between linguistically similar languages, such as Marathi and Konkani. The l3cube-pune/MarathiSentiment model particularly benefited from fine-tuning, as it showed stronger results for negative sentiments in Konkani. This suggests that, while fine-tuning on Marathi improves sentiment classification, incorporating Konkani-specific data in future experiments would further enhance accuracy across all sentiment classes, thereby contributing to more reliable NLP applications for low-resource languages.

The experiment was an attempt to show that although resource poor languages do not have datasets, they can be created with the help of similar languages.

The two models considered are as follows, Model 1: deepampatel/roberta-mlm-marathi and Model 2: l3cube-pune/MarathiSentiment.

5.1 Model 1 without Fine-Tuning (Marathi)

In the evaluation of Model 1 without fine-tuning, the overall accuracy achieved was 0.4233, with a balanced accuracy of 0.3112. Notably, the model performed well on negative sentiments, achieving an accuracy of 0.8623, while its performance on neutral and positive sentiments was significantly lower, at 0.0676 and 0.0037, respectively. This indicates that the model has a strong tendency to identify negative sentiments but struggles with neutral and positive classifications, as shown in Table 1 and Table 2.

Sentiment	Negative	Neutral	Positive
Negative	501	76	4
Neutral	69	5	0
Positive	485	58	2

Table 1: Confusion Matrix for Model 1 without Fine-Tuning (Marathi)

Metric	Value
Accuracy	0.4233
Balanced Accuracy	0.3112
Negative class Accuracy	0.8623
Neutral class Accuracy	0.0676
Positive class Accuracy	0.0037
F1 Macro	0.2222

Table 2: Accuracy Results for Model 1 without Fine-Tuning (Marathi

5.2 Model 1 with Fine-Tuning (Marathi)

Upon fine-tuning, Model 1 displayed a marginal increase in overall accuracy to 0.4275, with a balanced accuracy of 0.3136. The improvement in negative sentiment accuracy to 0.5370 suggests some capability in identifying negatives post-fine-tuning. However, the neutral and positive accuracy remained low at 0.0405 and 0.3633, respectively. These results indicate that while fine-tuning offers slight improvements, challenges in accurately predicting neutral and positive sentiments persist, as shown in Table 3 and Table 4.

5.3 Zero Shot: Analyzing Konkani with Marathi Sentiment on Model 1

For the zero-shot application of Model 1 on Konkani, the accuracy reached 0.4442, with a bal-

Sentiment	Negative	Neutral	Positive
Negative	312	12	257
Neutral	43	3	28
Positive	332	15	198

Table 3: Confusion Matrix for Model 1 with Fine-Tuning (Marathi)

Metric	Value
Accuracy	0.4275
Balanced Accuracy	0.3136
Negative class Accuracy	0.5370
Neutral class Accuracy	0.0405
Positive class Accuracy	0.3633
F1 Macro	0.3117

Table 4: Accuracy Results for Model 1 with Fine- Tuning (Marathi)

anced accuracy of 0.3282. The model's performance was particularly notable in classifying positive sentiments, achieving an impressive accuracy of 0.8936. However, it struggled with negative sentiments, reflected in an accuracy of only 0.0775. These findings highlight the model's strength in recognizing positive sentiments in Konkani, while its performance in identifying negative sentiments requires further enhancement, as shown in Table 5 and Table 6.

Sentiment	Negative	Neutral	Positive
Negative	45	1	535
Neutral	10	1	63
Positive	57	1	487

Table 5: Confusion Matrix for Model 1 on Konkani (Zero-Shot)

Metric	Value
Accuracy	0.4442
Balanced Accuracy	0.3282
Negative class Accuracy	0.0775
Neutral class Accuracy	0.0135
Positive class Accuracy	0.8936
F1 Macro	0.2511

Table 6: Accuracy Results for Model 1 on Konkani (Zero-Shot)

5.4 Model 2 without Fine-Tuning (Marathi)

Model 2's initial evaluation without fine-tuning resulted in an overall accuracy of 0.2450, with a balanced accuracy of 0.3725. The model displayed a modest ability to classify neutral sentiments accurately at 0.6892, while negative and positive accuracies were significantly lower at 0.2651 and 0.1633, respectively. This indicates that the model has a somewhat better grasp on neutral sentiments but generally performs poorly across the board, as shown in Table 7 and Table 8.

Sentiment	Negative	Neutral	Positive
Negative	154	354	73
Neutral	16	51	7
Positive	127	329	89

Table 7: Confusion Matrix for Model 2 without Fine-Tuning (Marathi)

Metric	Value
Accuracy	0.2450
Balanced Accuracy	0.3725
Negative class Accuracy	0.2651
Neutral class Accuracy	0.6892
Positive class Accuracy	0.1633
F1 Macro	0.2421

Table 8: Accuracy Results for Model 2 without Fine-Tuning (Marathi)

5.5 Model 2 with Fine-Tuning (Marathi)

After fine-tuning, Model 2 showed improvement with an overall accuracy of 0.4817 and a balanced accuracy of 0.3465. The model's negative sentiment accuracy improved to 0.6730, indicating a positive shift in performance for negative sentiment classification. However, the neutral accuracy dropped to 0.0270 and positive accuracy slightly improving to 0.3394. These metrics suggest that while fine-tuning enhanced the model's performance in recognizing negative sentiments, challenges remain in accurately predicting neutral sentiments, as shown in Table 9 and Table 10.

5.6 Zero-Shot: Analyzing Konkani with Marathi Sentiment on Model 2

The zero-shot application of Model 2 on Konkani yielded an accuracy of 0.4917 and a balanced accuracy of 0.3595. The model exhibited a solid performance in identifying positive sentiments, achieving

Sentiment	Negative	Neutral	Positive
Negative	391	12	178
Neutral	51	2	21
Positive	347	13	185

Table 9: Confusion Matrix for Model 2 with Fine-Tuning (Marathi)

Metric	Value
Accuracy	0.4817
Balanced Accuracy	0.3465
Negative Accuracy	0.6730
Neutral Accuracy	0.0270
Positive Accuracy	0.3394
F1 Macro	0.3362

Table 10: Accuracy Results for Model 2 with Fine-Tuning (Marathi)

an accuracy of 0.4440. However, its negative sentiment accuracy was moderate at 0.5938, and neutral sentiment accuracy was notably low at 0.0405. These results indicate that Model 2 effectively recognizes positive sentiments in Konkani, but its performance on negative and neutral sentiments still leaves room for improvement, as shown in Table 11 and Table 12.

Sentiment	Negative	Neutral	Positive
Negative	345	17	219
Neutral	40	3	31
Positive	296	7	242

Table 11: Confusion Matrix for Model 2 on Konkani (Zero-Shot)

Metric	Value
Accuracy	0.4917
Balanced Accuracy	0.3595
Negative Accuracy	0.5938
Neutral Accuracy	0.0405
Positive Accuracy	0.4440
Macro F1 Score	0.3576

Table 12: Accuracy Results for Model 2 on Konkani (Zero-Shot)

6 Conclusion

Based on the results presented, we observed that fine-tuning the models significantly improved sentiment classification accuracy for both Marathi and Konkani. Model 1 demonstrated better performance after fine-tuning, particularly in the zero-shot application to Konkani, where the positive sentiment classification was especially strong. Model 2 also performed well in the zero-shot learning scenario. The fine-tuned models showed better performance in sentiment classification across both languages and we expect that additional training on Konkani-specific data would further improve performance.

Our study highlights that the method used in this research can be employed to create sentiment-tagged data for Konkani despite the absence of large, annotated datasets. The cross-lingual transfer from Marathi to Konkani proves effective, making this approach a technique we look forward to for building sentiment-tagged datasets for low-resource languages like Konkani, where labeled data is limited.

References

Joseph Bloch. 1970. Formation of the Marathi Language. Motilal Banarsidass Publ.

A Conneau. 2019. Unsupervised cross-lingual representation learning at scale. arXiv preprint arXiv:1911.02116.

Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

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.

Mohammad Heydarian, Thomas E Doyle, and Reza Samavi. 2022. MLCM: Multi-label confusion matrix. IEEE Access, 10:19083–19095.

Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiusi Chen, and Julian McAuley. 2024. Bridging language and items for retrieval and recommendation. arXiv preprint arXiv:2403.03952.

Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. arXiv preprint arXiv:1801.06146.

P. Joshi, S. Santy, A. Budhiraja, K. Bali, and M. Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. arXiv preprint.

A. Kunchukuttan. 2020. The indicnlp library.

L3Cube-Pune. 2024. Marathi sentiment analysis model. https://huggingface.co/l3cube-pune/MarathiSentiment. Accessed: 2024-09-30.

Stella Mascarenhas-Keyes. 2020. Language and diaspora: The use of portuguese, english and konkani by catholic goan women. In Bilingual Women, pages 149–166. Routledge.

Walaa Medhat, Ahmed Hassan, and Hoda Korashy. 2014. Sentiment analysis algorithms and applications: A survey. Ain Shams Engineering Journal, 5(4):1093–1113.

Naeem Akhtar Muhammad, Abdul Rehman, and Usman Shoaib. 2017. Accuracy based feature ranking metric for multi-label text classification. International Journal of Advanced Computer Science and Applications, 8(10).

Deepam Patel. 2023. Roberta for masked language modeling in marathi. Accessed: 2024-09-30.

T Pires. 2019. How multilingual is multilingual bert. arXiv preprint arXiv:1906.01502.

M. Sardesai. 2006. A Comparative Linguistic and Cultural Study of Lexical Influences on Konkani. Doctoral dissertation, Goa University.