

## BLP-2023 Task 2: Sentiment Analysis

Md. Arid Hasan<sup>1</sup>, Firoj Alam<sup>2</sup>, Anika Anjum<sup>3</sup>, Shudipta Das<sup>3</sup>, Afiyat Anjum<sup>3</sup>

<sup>1</sup>SE+AI Research Lab, University of New Brunswick, Fredericton, Canada

<sup>2</sup>Qatar Computing Research Institute, Doha, Qatar

<sup>3</sup>Daffodil International University, Dhaka, Bangladesh  
arid.hasan@unb.ca, fialam@hbku.edu.qa

### Abstract

We present an overview of the BLP Sentiment Shared Task, organized as part of the inaugural BLP 2023 workshop, co-located with EMNLP 2023. The task is defined as the detection of sentiment in a given piece of social media text. This task attracted interest from 71 participants, among whom 29 and 30 teams submitted systems during the development and evaluation phases, respectively. In total, participants submitted 597 runs. However, a total of 15 teams submitted system description papers. The range of approaches in the submitted systems spans from classical machine learning models, fine-tuning pre-trained models, to leveraging Large Language Model (LLMs) in zero- and few-shot settings. In this paper, we provide a detailed account of the task setup, including dataset development and evaluation setup. Additionally, we provide a brief overview of the systems submitted by the participants. All datasets and evaluation scripts from the shared task have been made publicly available for the research community, to foster further research in this domain.<sup>1</sup>

### 1 Introduction

Sentiment analysis has emerged as a significant sub-field in Natural Language Processing (NLP), with a wide array of applications encompassing social media monitoring, brand reputation management, market research, customer feedback analysis, among others. The advancement of sentiment analysis systems has been driven by substantial research efforts, addressing its indispensable utility across diverse fields such as business, finance, politics, education, and services (Cui et al., 2023). Traditionally, analysis has been conducted across various types of content and domains including news articles, blog posts, customer reviews, and social media posts, and extended over different modali-

ties like textual and multimodal analyses (Hussein, 2018; Dashtipour et al., 2016).

At its core, the task of sentiment analysis is defined as the extraction and identification of polarities (e.g., positive, neutral, and negative) expressed within texts. However, its scope has broadened to encompass the identification of: (i) the target (i.e., an entity) or aspect of the entity on which sentiment is expressed, (ii) the opinion holder, and (iii) the time at which it is expressed (Liu, 2020). Such advancements have primarily been made for high-resource languages.

Research on fundamental sentiment analysis remains an ongoing exploration, especially for many low-resource languages, primarily due to the scarcity of datasets and consolidated community effort. Although there has been a recent surge in interest (Batanović et al., 2016; Nabil et al., 2015; Muhammad et al., 2023), the field continues to pose significant challenges. Similar to other low-resource languages, the challenges for sentiment analysis in Bangla have been reported in recent studies (Alam et al., 2021a; Islam et al., 2021, 2023). Alam et al. (2021a) emphasized the primary challenges associated with Bangla sentiment analysis, specifically issues of duplicate instances in the data, inadequate reporting of annotation agreement, and generalization. These challenges were also highlighted in (Islam et al., 2021), further emphasizing the need to address them for effective sentiment analysis in Bangla.

To advance research in Bangla sentiment analysis, we emphasized community engagement and organized a shared task at BLP 2023. Similar efforts have primarily been conducted for other languages as part of the SemEval Workshop. The analysis of sentiment in tweets serves as an example of such efforts, particularly focusing on Arabic and English (Rosenthal et al., 2017). An earlier attempt at such an endeavor for Bangla is reported in (Patra et al., 2015), which mainly focused on

<sup>1</sup>[https://github.com/blp-workshop/blp\\_task2](https://github.com/blp-workshop/blp_task2)

tweets. Our initiative significantly different from theirs in terms of datasets (e.g., data from multiple social media platforms and diverse domains) and evaluation setup.

A total of 71 teams registered for the task, out of which 30 made an official submission on the test set, and 15 of the participating teams submitted a system description paper.

The remainder of the paper is structured as follows: Section 2 provides an overview of the relevant literature. Section 3 discusses the task and dataset. Section 4 describes the organization of the task and the evaluation measures. An overview of the participating systems is provided in Section 5. Lastly, Section 6 concludes the paper.

## 2 Related Work

The current state-of-the-art research for Bangla sentiment classification mainly dominated focuses on two key aspects: the development of datasets and model development. Notable recent work in this direction include (Chowdhury and Chowdhury, 2014; Alam et al., 2021a; Islam et al., 2021; Kabir et al., 2023; Islam et al., 2023). Kabir et al. (2023) curated the largest dataset from book reviews, with annotations based on the review ratings. Although the dataset encompasses a large number of reviews, the class distribution poses a challenge for the Negative and Neutral classes. A well-balanced dataset has been explored in (Islam et al., 2021), comprising ~15K manually annotated comments spanning 13 different domains. This dataset is also used as a part of this shared task.

From a modeling perspective, the existing literature addresses the problem using both classical machine learning and deep learning algorithms. These include Naive Bayes, Support Vector Machine, Decision Tree, Maximum Entropy, and Random Forest (Rahman and Hossen, 2019; Banik and Rahman, 2018; Chowdhury et al., 2019; Islam et al., 2016). Moreover, recent studies have extensively employed deep learning models for Bangla sentiment classification (Hassan et al., 2016; Aziz Sharfuddin et al., 2018; Tripto and Ali, 2018; Ashik et al., 2019; Karim et al., 2020; Sazed, 2021; Sharmin and Chakma, 2021). Common deep learning approaches incorporate LSTMs, CNNs, attention mechanisms, and multichannel convolutional LSTMs. In the studies by Hasan et al. (2020); Alam et al. (2021a), comprehensive comparisons across various datasets were conducted, illustrating that

the deep learning-based pretrained language model XLM-RoBERTa excels in performance. Comparisons between classical and deep learning-based approaches have also been explored (Ashik et al., 2019; Hasan et al., 2020; Alam et al., 2021a).

Given the significant capabilities that Large Language Models (LLMs) have demonstrated across diverse applications and scenarios, Hasan et al. (2023) explored various LLMs such as Flan-T5 (large and XL) (Chung et al., 2022), Bloomz (1.7B, 3B, 7.1B, 176B-8bit) (Muennighoff et al., 2022), and GPT-4 (OpenAI, 2023), comparing the results with fine-tuned models. The resulting performance demonstrate that fine-tuned models continue to outperform zero- and few-shot prompting. However, the performance of LLMs showcases a promising direction towards the development of systems with limited datasets for new domains.

Though there is a surge of research interest and progress, utilizing such systems in real applications remains a challenge in terms of performance and generalization capability. This shared task aimed to advance research through community effort and focus on a standard evaluation setup. As a starting point, we aimed to classify sentiment into three sentiment polarities: positive, neutral, and negative. This approach can be further extended in future studies.

## 3 Task and Dataset

### 3.1 Task

The task is defined as “detect the sentiment associated within a given text”. This is a multi-class classification task that involves determining whether the sentiment expressed in the text is *Positive*, *Negative*, and *Neutral*.

### 3.2 Dataset

We utilized the MUBASE (Hasan et al., 2023) and SentNoB (Islam et al., 2021) datasets for the task. Both datasets were annotated by multiple annotators, with the inter-annotation agreement being 0.84 for MUBASE and 0.53 for SentNoB, respectively. The SentNoB data is curated from newspapers and YouTube video comments, covering 13 different topics such as Politics, National, International, Food, Sports, Teach, etc. The MUBASE dataset consists of comments from popular news media sources such as BBC Bangla, Prothom Alo, and BD24Live, which were collected from Facebook and Twitter.

We further analyzed the distribution of sentences based on the number of words associated with each class label, as depicted in Table 1. We created various ranges of sentence length buckets to understand and define the sequence length while training the transformer-based models. It appears that more than 80% of the posts comprise twenty words or fewer, a finding consistent with the typical of social media posts, as observed in previous studies (Alam et al., 2021b). Moreover, the average number of words and sentences per data point are 15.87 and 1.03, respectively.

Split	#Words	Pos	Neu	Neg
Train	<10	5,616	3,595	6,575
	11-20	4,587	2,212	5,613
	21-30	1,263	671	1,949
	31-40	493	287	818
	41-50	260	152	377
	51+	145	218	435
Dev	<10	587	398	723
	11-20	539	244	634
	21-30	160	68	232
	31-40	67	43	90
	41-50	22	14	34
	51+	13	26	40
Dev-test	<10	601	292	783
	11-20	420	178	603
	21-30	68	55	178
	31-40	11	21	54
	41-50	6	16	29
	51+	20	38	53
Test	<10	1,111	627	1,482
	11-20	762	382	1,183
	21-30	140	121	371
	31-40	31	56	111
	41-50	16	26	71
	51+	32	65	120

Table 1: Detailed class label distribution of the shared task data splits. Pos: Positive, Neu: Neutral, Neg: Negative.

Dataset	Train	Dev	DT	Test
MUBASE	✓	✗	✓	✓
SentNoB	✓	✓	✗	✗

Table 2: Data sources utilized in various splits for the shared task. DT: Dev-Test

For the shared task, we combined the MUBASE (Hasan et al., 2023) training set with the SentNoB (Islam et al., 2021) training set, resulting in a total of 35,266 entries for the training set. The SentNoB development set was used as the shared task development set. Additionally, the MUBASE development set served as the dev-test set for the shared task, while the test set was utilized for system evaluation and participant ranking. The specifics of the data sources are outlined in Table 2, and the detailed distribution of the data split is presented in Table 3.

Class	Train	Dev	DT	Test	Total
Pos	12,364	1,388	1,126	2,092	16,970
Neu	7,135	793	600	1,277	9,805
Neg	15,767	1,753	1,700	3,338	22,558
<b>Total</b>	<b>35,266</b>	<b>3,934</b>	<b>3,426</b>	<b>6,707</b>	<b>49,333</b>

Table 3: Class label distribution of the shared task dataset. DT: Dev-Test, Pos: Positive, Neu: Neutral, Neg: Negative

## 4 Evaluation Framework

### 4.1 Evaluation Measures

For evaluation, we used the *Micro-F1 score* and the evaluation scripts along with data are available online<sup>2</sup>. As reference points, we provided both the majority and random baselines. The majority baseline always predicts the most common class in the training data and assigns this class to each instance in the test dataset. Conversely, the random baseline assigns one of the classes randomly to each instance in the test dataset.

### 4.2 Task Organization

For the shared task, we provided four sets of data: the training set, development set, development-test set, and test set, as outlined in Table 3. The purpose of providing the development set is for hyperparameter tuning. We provided the development test set without labels to allow participants to evaluate their systems during the system development phase. The test set was designated for the final system evaluation and ranking. We ran the shared task in two phases and hosted the submission system on the CodaLab platform.<sup>3</sup>

<sup>2</sup>[https://github.com/blp-workshop/blp\\_task2](https://github.com/blp-workshop/blp_task2)

<sup>3</sup><https://codalab.lisn.upsaclay.fr/competitions/14587>

**Development Phase** In the first phase, only the training set, development set, and development-test set were made available, with no gold labels provided for the latter. Participants competed against each other to achieve the best performance on the development test set. A live leaderboard was made available to keep track of all submissions.

**Test Phase** In the second phase, the test set was released without labels, and the participants were given just four days to submit their final predictions. The test set was used for evaluation and ranking. The leaderboard was set to private during the evaluation phase, and participants were allowed to submit multiple systems without seeing the scores. The last valid submission was considered for official ranking.

After the competition concluded, we released the test set with gold labels to enable participants to conduct further experiments and error analysis.

## 5 Results and Overview of the Systems

### 5.1 Results

A total of 29 and 30 teams submitted their systems during the development and evaluation phases, respectively. In Table 4, we report the results of the submitted system on dev-test and test sets. We also include the results for the majority and random baselines. The ranking on the table was determined by the results from the test set. Note that some teams participated in the development phase but did not participate in the evaluation phase, and vice versa, as indicated by the symbol  $\times$ . Additionally, the team marked with \* did not submit a system description paper.

Upon comparing the results from the dev-test and test sets across different teams, it appears that the performance difference between them is very minimal. The models did not exhibit overfitting; in some cases, the performance on the test set even surpassed that on the dev-test set.

As can be seen in Table 4, almost all systems outperformed random baseline except one system, whereas 26 systems outperformed the majority baseline. The best system, Aambela (Fahim, 2023), achieved micro-F1 score of 0.73, which is an absolute improvement of 0.23. The team mainly fine-tuned BanglaBERT and multilingual BERT along with adversarial weight perturbation. The second best system, Knowdee (Liu et al., 2023), used data augmentation with psudolabeling, which are ob-

tained from an ensemble of models. The third best system, LowResource (Chakma and Hasan, 2023), used ensemble of different fine-tuned models.

In Table 5, we report the overview of the approaches of the submitted systems. The most used models are multilingual BERT, BanglaBERT, and XLM-RoBERTa. Specifically, 9, 8, and 14 out of 15 teams utilized multilingual BERT, BanglaBERT, and XLM-RoBERTa, respectively. Ensembles of fine-tuned models provide the best systems for this task. Additionally, two teams applied few-shot learning using the mT5, BanglaBERT large, and GPT-3.5 models. However, the teams did not provide the details regarding the prompts.

### 5.2 Discussion

From the official ranking presented in Table 4, early every team outperformed the performance of the random baseline system. The performance difference between the top 22 teams is very small compared with the 23rd-ranked team. In Table 6, we presented the per-class performances for the top 5 teams. Although most of the teams performed better than the random baseline by a large margin, the neutral class is still the most difficult one to identify. The low performance in neutral class might be due to its skewed distribution in the dataset. Data augmentation, up-sampling the minority class, and class re-weighting are common approaches typically used to address such issues. Although some systems employed data augmentation, it seems this issue was not thoroughly considered across all teams.

### 5.3 Participating Systems

Below, we provide a brief description of the participating systems and their leaderboard rank.

**Aambela (Fahim, 2023)** (rank 1) emerged as the best-performing team in the shared task, fine-tuning pretrained models BanglaBERT (Bhattacharjee et al., 2022a) and multilingual BERT (Devlin et al., 2019) using two classification heads. Initially, the author removed URLs and HTML tags, then applied a normalizer to the preprocessed text. Adversarial weight perturbation was utilized to enhance the training’s robustness, and a 5-fold cross-validation was also conducted.

**Knowdee (Liu et al., 2023)** (rank 2) partitioned the data set into 10 folds and generated pseudo-labels for unlabeled data using a fine-tuned ensemble of models. They employed standard data

Rank – Team	Micro-F1	
	Dev-Test	Test
1. Aambela (Fahim, 2023)	0.7303	0.7310
2. Knowdee (Liu et al., 2023)	0.7288	0.7267
3. LowResource (Chakma and Hasan, 2023)	0.7224	0.7179
4. LowResourceNLU (Veeramani et al., 2023)	0.7248	0.7172
5. Z-Index (Tarannum et al., 2023)	✗	0.7164
- ShadmanRohan*	0.7207	0.7155
6. RGB*	0.7125	0.7112
7. EmptyMind(Fatema et al., 2023)	0.7215	0.7109
8. KeAb*	0.7125	0.7094
9. Embeddings (Tonmoy, 2023)	✗	0.7088
10. RSM-NLP (Seth et al., 2023)	0.7023	0.7078
11. DeepBlueAI*	✗	0.7076
12. nlpBDpatriots (Goswami et al., 2023)	0.7192	0.7058
13. NLP_CUET*	0.6278	0.7052
14. M1437 (Rahman and Uzuner, 2023)	0.7315	0.7036
15. Semantic_Savants*	0.6961	0.7002
16. meemaw*	✗	0.6996
17. Score_IsAll_You_Need*	0.6909	0.6930
18. VishwasGPai*	0.6970	0.6824
19. UFAL-ULD (Mukherjee et al., 2023)	0.6661	0.6768
20. Semantics Squad (Dey et al., 2023)	0.7201	0.6742
21. BanglaNLP (Saha and Nanda, 2023)	0.6810	0.6702
22. VacLM*	✗	0.6584
23. trina*	✗	0.6194
- Rachana8_K*	✗	0.5962
24. lixn*	✗	0.5889
25. Baseline (Majority)	0.4962	0.4977
26. Xenon*	✗	0.4534
27. Error Point (Das et al., 2023)	✗	0.4129
28. SSCP*	0.5584	0.3390
29. Baseline (Random)	0.3389	0.3356
30. Ushoshi2023 (Khushbu et al., 2023)	✗	0.2626
- Shilpa*	0.7166	✗
- Dhiman*	0.7154	✗
- KarbonDark*	0.7154	✗
- MrinmoyMahato*	0.7107	✗
- shakib034*	0.6734	✗
- Saumajit*	0.6559	✗
- sankalok*	0.6203	✗
- DiscoDancer420*	0.5736	✗
- Devs*	0.5736	✗
- almamunsardar*	0.5642	✗

Table 4: Official ranking of the shared task on the test set. \*No working note submitted. - Run submitted after the deadline. ✗ - indicates team has not submitted system in the respective phase.

preprocessing and augmentation techniques to process the data, and fine-tuned BanglaBERT (Bhattacharjee et al., 2022a), MuRIL (Khanuja et al., 2021), XLM-RoBERTa (Conneau et al., 2020), and

mT5 (Xue et al., 2021), achieving the second-best performance. The team also implemented Few-shot (3-shot) learning and compared the results with those from fine-tuned models.

Team	Models														Misc.			
	Classical	multilingual BERT	RoBERTa	XLNet	BanglaBERT	BanglishBERT	MuRIL	mT5	BanglaT5	Indic-BERT	BanglaGPT2	DistilBERT	LSTM	LSTM-CNN	Bangla-BERT	Few-shot	Ensemble	Preprocessing
Aambela (Fahim, 2023)		✓			✓												✓	✓
Knowdee (Liu et al., 2023)				✓	✓		✓	✓									✓	✓
LowResource (Chakma and Hasan, 2023)				✓	✓			✓									✓	✓
LowResourceNLU (Veeramani et al., 2023)		✓			✓												✓	✓
Z-Index (Tarannum et al., 2023)	✓	✓			✓													✓
EmptyMind (Fatema et al., 2023)	✓				✓							✓						✓
Embeddings (Tonmoy, 2023)		✓			✓				✓	✓								✓
RSM-NLP (Seth et al., 2023)			✓		✓	✓	✓							✓		✓	✓	✓
nlpBDpatriots (Goswami et al., 2023)	✓	✓		✓	✓		✓									✓		✓
M1437 (Rahman and Uzuner, 2023)				✓	✓													✓
UFAL-ULD (Mukherjee et al., 2023)		✓		✓	✓									✓			✓	✓
Semantics Squad (Dey et al., 2023)		✓		✓	✓		✓										✓	✓
BanglaNLP (Saha and Nanda, 2023)	✓	✓		✓	✓												✓	✓
Error Point (Das et al., 2023)	✓											✓	✓				✓	✓
Ushoshi2023 (Khushbu et al., 2023)	✓	✓		✓	✓						✓	✓					✓	✓

Table 5: Overview of the approaches used in the submitted systems.

**LowResource (Chakma and Hasan, 2023)** (rank 3) fine-tuned both the base and large versions of BanglaBERT (Bhattacharjee et al., 2022a), employing randomly dropping tokens, and also fine-tuned XLM-RoBERTa (Conneau et al., 2020). During the development phase, they created an ensemble of three models. However, for the evaluation phase, they ensembled only two variants of BanglaBERT, with one of them being fine-tuned using external data. Additionally, they employed task-adaptive pretraining and paraphrasing techniques utilizing BanglaT5 (Bhattacharjee et al., 2022b).

**LowResourceNLU (Veeramani et al., 2023)** (rank 4) fine-tuned BanglaBERT base and large (Bhattacharjee et al., 2022a), with MLM and classification heads, and multilingual BERT (Devlin et al., 2019) jointly on the XNLI and shared task dataset. They also created an ensemble of all three transformer-based models and applied multi-step aggregation to capture the most confident class predicted across all models.

**Z-Index (Tarannum et al., 2023)** (rank 5) utilized standard preprocessing techniques to remove URLs, usernames, emojis, and hashtags from the text. Initially, they employed SVM and Random Forest classical models, and later fine-tuned both the base and large variants of BanglaBERT (Bhattacharjee et al., 2022a), as well as the multilingual BERT (Devlin et al., 2019). The model was trained using the provided training set.

**EmptyMind (Fatema et al., 2023)** (rank 7) initially applied classical models such as Decision Tree, Random Forest, SVM, and XGBoost, utilizing TF-IDF vectors, as well as Word2Vec vectors. Subsequently, they employed deep learning-based models including Stacked BiLSTM and BiLSTM+CNN. Furthermore, they fine-tuned different variants of BanglaBERT (Bhattacharjee et al., 2022a).

**Embeddings (Tonmoy, 2023)** (rank 9) fine-tuned pretrained models BanglaBERT (Bhattacharjee et al., 2022a), BanglaGPT2,<sup>4</sup> Indic-BERT (Kakwani et al., 2020), and multilingual BERT (Devlin et al., 2019) using cross entropy loss function. Later to reduce the computational cost, they investigated the performances across the self-adjusting dice loss, focal loss, and F1-micro loss. They also combined training, dev, and dev-test sets as training data to train and test data to evaluate the performances of the models.

**RSM-NLP (Seth et al., 2023)** (rank 10) submitted their runs by fine-tuning RoBERTa (Liu et al., 2019), DistilBERT (Sanh et al., 2019), BanglaBERT,<sup>5</sup> BanglaBERT (Bhattacharjee et al., 2022a), BanglishBERT (Bhattacharjee et al., 2022a), and MuRIL (Khanuja et al., 2021), with the additional use of training data. They employed standard pre-

<sup>4</sup><https://huggingface.co/flax-community/gpt2-bengali>

<sup>5</sup><https://github.com/sagorbrur/bangla-bert>

Class	Baseline	Aambela	Knowdee	LowResource	LowResourceNLU	Z-Index
Negative	0.3996	0.7958	0.7943	0.7873	0.7877	0.7877
Neutral	0.2368	0.4998	0.4592	0.3998	0.4021	0.4250
Positive	0.3329	0.7666	0.7599	0.7567	0.7530	0.7559

Table 6: F1 scores of the baseline and top five systems for each class.

processing techniques to process the data. They also submitted ensemble techniques (i.e., weighted and majority-voted) of fine-tuned models.

**nlpBDpatriots (Goswami et al., 2023)** (rank 12) began with traditional approaches such as logistic regression and SVM. Later, they fine-tuned BanglaBERT (Bhattacharjee et al., 2022a), multilingual BERT (Devlin et al., 2019), MuRIL (Khanuja et al., 2021), and XLM-RoBERTa (Conneau et al., 2020), and ensemble the models using a weighted average of the confidence predicted by each model. They also employed few-shot learning using GPT-3.5 (OpenAI, 2023).

**M1437 (Rahman and Uzuner, 2023)** (rank 14) fine-tuned large pretrained language models BanglaBERT large (Bhattacharjee et al., 2022a) and XLM-RoBERTa large (Conneau et al., 2020) along with the base version of each model. They also used an existing dataset (Hasan et al., 2020) in addition to the provided training data. To compare among the transformers models, they also fine-tuned the multilingual BERT. During the development phase, they were the best-performing team and they ended the competition in the 14th position in the evaluation phase.

**UFAL-ULD (Mukherjee et al., 2023)** (rank 19) fine-tuned BanglaBERT (Bhattacharjee et al., 2022a), Bangla-BERT<sup>6</sup> multilingual BERT (Devlin et al., 2019), and XLM-RoBERTa (Conneau et al., 2020) to tackle the problem. They followed the standard preprocessing steps to process the data and upsampled the training data to achieve balance among the classes. They also employed a focal loss function to address hard-to-classify examples.

**Semantics Squad (Dey et al., 2023)** (rank 20) submitted runs for both the development and evaluation phases. Standard preprocessing techniques were applied, with URLs and hashtags being removed from the data, to process and fine-tune BanglaBERT (Bhattacharjee et al., 2022a), BanglaBERT (Bhattacharjee et al., 2022a), XLM-

RoBERTa (Conneau et al., 2020), and multilingual BERT (Devlin et al., 2019).

**BanglaNLP (Saha and Nanda, 2023)** (rank 21) also fine-tuned BanglaBERT (Bhattacharjee et al., 2022a), BERT multilingual (Devlin et al., 2019), and XLM-RoBERTa (Conneau et al., 2020) pre-trained models. Additionally, they performed parameter-efficient tuning (P-tuning) on XLM-RoBERTa. They also employed traditional models such as Logistic Regression, Naive Bayes, SGD Classifier, Majority Voting, and Stacking in their approach to the task.

**Error Point (Das et al., 2023)** (rank 27) performed preprocessing by removing duplicate text, filtering based on text length, and eliminating punctuation, links, emojis, non-character elements, and stopwords. They also carried out data augmentation. For their analysis, they utilized classical algorithms such as Logistic Regression, Decision Tree, Random Forest, Multinomial Naive Bayes, SVM, and SGD, using  $n$ -grams to represent the input. Additionally, they employed deep learning models, namely LSTM and LSTM-CNN.

**Ushoshi2023 (Khushbu et al., 2023)** (rank 30) applied preprocessing by removing punctuation marks, links, emojis, hashtag signs, usernames, and non-Bangla characters. They also applied an upsampling technique to balance the dataset. Initially, they employed traditional models such as logistic regression, decision tree, random forest, multinomial naive bayes, k-nearest neighbor, SVM, and SGD for classification. Subsequently, they fine-tuned BanglaBERT (Bhattacharjee et al., 2022a), XLM-RoBERTa (Conneau et al., 2020), DistilBERT (Sanh et al., 2019), and multilingual BERT (Devlin et al., 2019). Additionally, they trained a deep learning model, LSTM, to compare the performances across different models.

## 6 Conclusion and Future Work

We presented an overview of the shared task 2 (sentiment analysis) at the BLP Workshop 2023. Task 2 aimed to classify the sentiment in textual

<sup>6</sup><https://github.com/sagorbrur/bangla-bert>

content. Notable systems employed an ensemble of pretrained language models, with the language-specific BanglaBERT being the most popular. Also, some interesting approaches including P-tuning, Few-shot learning, LLMs, and different loss functions have been explored for tackling the problem. In general, numerous models, including different kinds of transformers, have been used in the current submissions for the task.

In future work, we plan to extend the task in various ways, such as aspect-based sentiment analysis and incorporating multiple modalities.

## Limitations

The BLP-2023 sentiment analysis shared task primarily focuses on sentiment polarity classification (positive, negative, and neutral) at the post level. This approach limits the identification of specific sentiment aspects and other crucial elements associated with them. Future editions of the task will address this aspect. Moreover, this edition focused solely on unimodality (text-only) models, leaving multimodal models for future study.

## References

- Firoj Alam, Md. Arid Hasan, Tanvirul Alam, Akib Khan, Jannatul Tajrin, Naira Khan, and Shammur Absar Chowdhury. 2021a. A review of bangla natural language processing tasks and the utility of transformer models. *arXiv preprint arXiv:2107.03844*.
- Firoj Alam, Umair Qazi, Muhammad Imran, and Ferda Offi. 2021b. Humaid: Human-annotated disaster incidents data from twitter with deep learning benchmarks. In *Proceedings of the International AAAI Conference on Web and social media*, volume 15, pages 933–942.
- Md Akhter-Uz-Zaman Ashik, Shahriar Shovon, and Summit Haque. 2019. Data set for sentiment analysis on bengali news comments and its baseline evaluation. In *Proc. of ICBSLP*, pages 1–5. IEEE.
- Abdullah Aziz Sharfuddin, Md. Nafis Tihami, and Md. Saiful Islam. 2018. A deep recurrent neural network with bilstm model for sentiment classification. In *2018 International Conference on Bangla Speech and Language Processing (ICBSLP)*, pages 1–4. IEEE.
- Nayan Banik and Md Hasan Hafizur Rahman. 2018. Evaluation of naïve bayes and support vector machines on Bangla textual movie reviews. In *Proc. of ICBSLP*, pages 1–6. IEEE.
- Vuk Batanović, Boško Nikolić, and Milan Milosavljević. 2016. Reliable baselines for sentiment analysis in resource-limited languages: The serbian movie review dataset. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 2688–2696.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022a. BanglaBERT: Language model pretraining and benchmarks for low-resource language understanding evaluation in Bangla. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1318–1327, Seattle, United States. Association for Computational Linguistics.
- Abhik Bhattacharjee, Tahmid Hasan, Wasi Uddin Ahmad, and Rifat Shahriyar. 2022b. Banglanlg: Benchmarks and resources for evaluating low-resource natural language generation in bangla. *CoRR*, abs/2205.11081.
- Aunabil Chakma and Masum Hasan. 2023. Low-resource at BLP-2023 Task 2: Leveraging banglabert for low resource sentiment analysis of bangla language. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Rumman Rashid Chowdhury, Mohammad Shahadat Hossain, Sazzad Hossain, and Karl Andersson. 2019. Analyzing sentiment of movie reviews in Bangla by applying machine learning techniques. In *Proc. of (ICBSLP)*, pages 1–6. IEEE.
- Shaika Chowdhury and Wasifa Chowdhury. 2014. Performing sentiment analysis in Bangla microblog posts. In *2014 International Conference on Informatics, Electronics Vision (ICIEV)*, pages 1–6.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 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, ACL '20*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jingfeng Cui, Zhaoxia Wang, Seng-Beng Ho, and Erik Cambria. 2023. Survey on sentiment analysis: evolution of research methods and topics. *Artificial Intelligence Review*, pages 1–42.
- Rajesh Kumar Das, Kabid Yeiad, Jannatul Maowa, Moshfiqur Rahman Ajmain, and Mirajul Islam. 2023. Team error point at BLP-2023 Task 2: A comparative exploration of hybrid deep learning and machine learning approach for advanced sentiment analysis



- techniques. In *Proceedings of the 1st Workshop on Bangla Language Processing (BLP 2023)*, Singapore. Association for Computational Linguistics.
- Kia Dashtipour, Soujanya Poria, Amir Hussain, Erik Cambria, Ahmad YA Hawalah, Alexander Gelbukh, and Qiang Zhou. 2016. Multilingual sentiment analysis: state of the art and independent comparison of techniques. *Cognitive computation*, 8:757–771.
- 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, NAACL-HLT '19*, pages 4171–4186, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Krishno Dey, Md. Arid Hasan, Prerona Tarannum, and Francis Palma. 2023. Semantics squad at BLP-2023 Task 2: Sentiment analysis of bengali text with fine tuned transformer based models. In *Proceedings of the 1st Workshop on Bangla Language Processing (BLP 2023)*, Singapore. Association for Computational Linguistics.
- Md Fahim. 2023. Aambela at BLP-2023 Task 2: Enhancing banglabert performance for bangla sentiment analysis task with in task pretraining and adversarial weight perturbation. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Karnis Fatema, Udoy Das, Md Ayon Mia, Md Sajidul Mowla, Mahshar Yahan, MD Fayeze Ullah, Arpita Sarker, and Hasan Murad. 2023. Emptymind at BLP-2023 Task 2: Sentiment analysis of bangla social media posts using transformer-based models. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Dhiman Goswami, Md Nishat Raihan, and Sadiya Sayara Chowdhury Puspo. 2023. nlpbdpatriots at BLP-2023 Task 2: A transfer learning approach towards bangla sentiment analysis. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Md Arid Hasan, Shudipta Das, Afiyat Anjum, Firoj Alam, Anika Anjum, Avijit Sarker, and Sheak Rashed Haider Noori. 2023. Zero-and few-shot prompting with llms: A comparative study with finetuned models for bangla sentiment analysis. *arXiv preprint arXiv:2308.10783*.
- Md Arid Hasan, Jannatul Tajrin, Shammur Absar Chowdhury, and Firoj Alam. 2020. Sentiment classification in bangla textual content: a comparative study. In *2020 23rd International Conference on Computer and Information Technology (ICCIT)*, pages 1–6. IEEE.
- Asif Hassan, Mohammad Rashedul Amin, Abul Kalam Al Azad, and Nabeel Mohammed. 2016. Sentiment analysis on Bangla and romanized Bangla text using deep recurrent models. In *2016 International Workshop on Computational Intelligence (IWCI)*, pages 51–56. IEEE.
- Doaa Mohey El-Din Mohamed Hussein. 2018. A survey on sentiment analysis challenges. *Journal of King Saud University-Engineering Sciences*, 30(4):330–338.
- Khondoker Ittehadul Islam, Sudipta Kar, Md Saiful Islam, and Mohammad Ruhul Amin. 2021. SentNoB: A dataset for analysing sentiment on noisy bangla texts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3265–3271.
- Md Ekramul Islam, Labib Chowdhury, Faisal Ahamed Khan, Shazzad Hossain, Sourave Hossain, Mohammad Mamun Or Rashid, Nabeel Mohammed, and Mohammad Ruhul Amin. 2023. SentiGOLD: A large bangla gold standard multi-domain sentiment analysis dataset and its evaluation. *arXiv preprint arXiv:2306.06147*.
- Md Saiful Islam, Md Ashiqul Islam, Md Afjal Hossain, and Jagoth Jyoti Dey. 2016. Supervised approach of sentimentality extraction from Bengali facebook status. In *2016 19th International Conference on Computer and Information Technology (ICCIT)*, pages 383–387. IEEE.
- Mohsinul Kabir, Obayed Bin Mahfuz, Syed Rifat Raiyan, Hasan Mahmud, and Md Kamrul Hasan. 2023. [BanglaBook: A large-scale bangla dataset for sentiment analysis from book reviews](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1237–1247, Toronto, Canada. Association for Computational Linguistics.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, NC Gokul, Avik Bhattacharyya, Mitesh M Khapra, and Pratyush Kumar. 2020. IndicNLPsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961.
- Md. Rezaul Karim, Bharathi Raja Chakravarthi, John P. McCrae, and Michael Cochez. 2020. [Classification benchmarks for under-resourced Bengali language based on multichannel convolutional-lstm network](#). *CoRR*, abs / 2004.07807.
- Simran Khanuja, Diksha Bansal, Sarvesh Mehtani, Savya Khosla, Atreyee Dey, Balaji Gopalan, Dilip Kumar Margam, Pooja Aggarwal, Rajiv Teja Nagipogu, Shachi Dave, et al. 2021. MuriL: Multilingual representations for indian languages. *arXiv preprint arXiv:2103.10730*.
- Sharun Akter Khushbu, Nasheen Nur, Mohiuddin Ahmed, and Nashtarin Nur. 2023. Ushoshi2023 at

- BLP-2023 Task 2: A comparison of traditional to advanced linguistic models to analyze sentiment in bangla texts. In *Proceedings of the 1st Workshop on Bangla Language Processing (BLP 2023)*, Singapore. Association for Computational Linguistics.
- Bing Liu. 2020. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge university press.
- Xiaoyi Liu, Teng Mao, Shuangtao Yang, and Bo Fu. 2023. Knowdee at BLP-2023 Task 2: Improving bangla sentiment analysis using ensembled models with pseudo-labeling. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. 2022. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786*.
- Shamsuddeen Hassan Muhammad, Idris Abdulmumin, Abinew Ali Ayele, Nedjma Ousidhoum, David Ifeoluwa Adelani, Seid Muhie Yimam, Ibrahim Sa'id Ahmad, Meriem Beloucif, Saif Mohammad, Sebastian Ruder, et al. 2023. Afrisenti: A twitter sentiment analysis benchmark for african languages. *arXiv preprint arXiv:2302.08956*.
- Sourabrata Mukherjee, Atul Kr. Ojha, and Ondřej Dušek. 2023. Ufal-uld at blp-2023 task 2 sentiment classification in bangla text. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Mahmoud Nabil, Mohamed Aly, and Amir Atiya. 2015. Astd: Arabic sentiment tweets dataset. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, pages 2515–2519.
- R OpenAI. 2023. Gpt-4 technical report. *arXiv*, pages 2303–08774.
- Braja Gopal Patra, Dipankar Das, Amitava Das, and Rajendra Prasath. 2015. Shared task on sentiment analysis in indian languages (sail) tweets-an overview. In *Mining Intelligence and Knowledge Exploration: Third International Conference, MIKE 2015, Hyderabad, India, December 9-11, 2015, Proceedings 3*, pages 650–655. Springer.
- A. Rahman and M. S. Hossen. 2019. Sentiment analysis on movie review data using machine learning approach. In *2019 International Conference on Bangla Speech and Language Processing (ICBSLP)*, pages 1–4.
- Majidur Rahman and Özlem Uzuner. 2023. M1437 at BLP-2023 Task 2: Harnessing bangla text for sentiment analysis: A transformer-based approach. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 502–518, Vancouver, Canada. Association for Computational Linguistics.
- Saumajit Saha and Albert Nanda. 2023. Banglanlp at BLP-2023 Task 2: Benchmarking different transformer models for sentiment analysis of bangla social media posts. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Salim Sazzed. 2021. Improving sentiment classification in low-resource bengali language utilizing cross-lingual self-supervised learning. In *International Conference on Applications of Natural Language to Information Systems*, pages 218–230. Springer.
- Pratinav Seth, Rashi Goel, Komal Mathur, and Swetha Vemulapalli. 2023. Rsm-nlp at BLP-2023 Task 2: Bangla sentiment analysis using weighted and majority voted fine-tuned transformers. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Sadia Sharmin and Danial Chakma. 2021. Attention-based convolutional neural network for bangla sentiment analysis. *Ai & Society*, 36(1):381–396.
- Prerona Tarannum, Md. Arid Hasan, Krishno Dey, and Sheak Rashed Haider Noori. 2023. Z-index at BLP-2023 Task 2: A comparative study on sentiment analysis. In *Proceedings of the 1st Workshop on Bangla Language Processing (BLP 2023)*, Singapore. Association for Computational Linguistics.
- S.M Towhidul Islam Tonmoy. 2023. Embeddings at BLP-2023 Task 2: Optimizing fine-tuned transformers with cost-sensitive learning for multiclass sentiment analysis. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.
- Nafis Irtiza Tripto and Mohammed Eunos Ali. 2018. Detecting multilabel sentiment and emotions from Bangla youtube comments. In *2018 International Conference on Bangla Speech and Language Processing (ICBSLP)*, pages 1–6. IEEE.

Hariram Veeramani, Surendrabikram Thapa, and Usman Naseem. 2023. Lowresourcenlu at blp: Enhancing sentiment classification and violence incitement detection through aggregated language models. In *Proceedings of the 1st International Workshop on Bangla Language Processing (BLP-2023)*, Singapore. Association for Computational Linguistics.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mt5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498.