

# Zero-Shot Multi-Label Classification of Bangla Documents: Large Decoders Vs. Classic Encoders

Souvika Sarkar<sup>1</sup>, Md. Najib Hasan<sup>1</sup>, Santu Karmaker<sup>2</sup>

<sup>1</sup> Accessible AI Lab (A2I Lab), School of Computing, Wichita State University

<sup>2</sup> Bridge-AI Lab, Department of Computer Science, University of Central Florida  
souvika.sarkar@wichita.edu, mxhasan39@shockers.wichita.edu, santu@ucf.edu

## Abstract

Bangla, a language spoken by over 300 million native speakers and ranked as the sixth most spoken language worldwide, presents unique challenges in natural language processing (NLP) due to its complex morphological characteristics and limited resources. Although recent large-decoder-based LLMs, such as GPT, LLaMA, and DeepSeek, have demonstrated excellent performance across many NLP tasks, their effectiveness in Bangla remains largely unexplored. In this paper, we establish the first benchmark comparing large decoder-based LLMs with classic encoder-based models for the Zero-Shot Multi-Label Classification (Zero-Shot-MLC) task in Bangla. Our evaluation of 32 state-of-the-art models reveals that existing so-called powerful encoders and decoders still struggle to achieve high accuracy on the Bangla Zero-Shot-MLC task, suggesting a need for more research and resources for Bangla NLP.

## 1 Introduction

Bangla is the sixth most spoken language worldwide, with over 300 million native speakers, accounting for nearly 4% of the global population<sup>1</sup>. Despite its widespread use, Bangla remains a low-resource language in the domain of Natural Language Processing (NLP) (Joshi et al., 2020), with limited benchmark studies and scarce linguistic resources. The rise of classic encoders, such as LASER (Artetxe and Schwenk, 2019) and LaBSE (Feng et al., 2020), alongside large decoder-based models such as GPT (Brown et al., 2020), BLOOM (Scao et al., 2022), LLaMA (Touvron et al., 2023), and DeepSeek (DeepSeek-AI et al., 2025), has driven significant advancement in multilingual NLP. However, while these models have demonstrated exceptional performance

in high-resource languages such as English, Chinese, and Spanish, their effectiveness in low-resource languages, like Bangla, remains largely unexplored. This gap raises a crucial question: How well do LLMs perform on Bangla-specific tasks compared to classic encoder-based models?

To address this, we present a benchmarking study, evaluating 32 state-of-the-art models on the Zero-Shot Multi-Label Classification (Zero-Shot-MLC) task for Bangla, including three classic encoders—LaBSE (Feng et al., 2020), LASER (Artetxe and Schwenk, 2019), and BanglaTransformers (Muhammad Rafsan Kabir et al., 2024)—and 29 large decoder-based LLMs. Our findings reveal that while a few models achieve  $F_1$  scores around 60%, most struggle with reliable classification, highlighting significant gaps in multilingual adaptability. This study provides the first systematic decoder vs. encoder evaluation for Bangla Zero-Shot-MLC, emphasizing the need for further research and resources in Bangla NLP. Our dataset and source code are publicly available at: <https://github.com/MdNajibHasan/BanglaNLP.git>.

## 2 Background and Related Works

**Bangla NLP and Zero-shot classification:** In the domain of Bangla NLP, several tasks have been explored by researchers in the past, including Information Extraction (Rahman et al., 2008; Uddin et al., 2019; Sharif et al., 2016), Machine Translation (Hasan et al., 2019; Anwar et al., 2009; Ismail et al., 2014), Named Entity Recognition (Banik and Rahman, 2018; Chaudhuri and Bhattacharya, 2008), Question Answering (Islam et al., 2016; Sarker et al., 2019; Kowsher et al., 2019), Sentiment Analysis (Uddin et al., 2019; Hassan et al., 2016; Alam et al., 2017), Spam and Fake Detection (Islam et al., 2019; Hussain et al., 2020) Annotator Bias Detection (Das et al., 2024), Hallucination Detection (Das et al., 2025) etc. Recently,

<sup>1</sup>Source: *List of languages by the number of native speakers*, Wikipedia, [https://en.wikipedia.org/wiki/List\\_of\\_languages\\_by\\_number\\_of\\_native\\_speakers](https://en.wikipedia.org/wiki/List_of_languages_by_number_of_native_speakers) (accessed on 23 May 2023)

Kowsher et al. (2022) introduce Bangla-BERT, a monolingual BERT model for the Bangla language. Bhattacharjee et al. (2022) presented BanglaBERT – a BERT-based (Devlin et al., 2019) Bangla NLU model pre-trained on 27.5GB data. Transformer-based models have been frequently employed in Bangla NLP research, such as in abusive comment detection (Aurpa et al., 2022b), text classification (Alam et al., 2020), question answering (Aurpa et al., 2022a; Adnan and Anwar, 2022), machine translation (Dhar et al., 2022), image caption generation (Palash et al., 2022), etc. Researchers have also introduced different datasets such as BanglaNLG (Bhattacharjee et al., 2023), a comprehensive benchmark for evaluating natural language generation models in Bangla, a sentiment analysis dataset SentNoB (Islam et al., 2021), a dataset for Bangla fake news detection (Hossain et al., 2020).

**Large Decoder (LLM) Based Generative Classification:** Recent research has extensively studied the potential of large language models like ChatGPT, Flan, BLOOM, etc., for a wide range of applications. For example, researchers have shown the utility of ChatGPT in healthcare education (Sallam, 2023), programming bug solving (Surameery and Shakor, 2023), and machine translation (Jiao et al., 2023).

**Classic Encoder Based Discriminative Classification:** Various topic modeling-based approaches have been explored for zero-shot classification for the English language (Karmaker Santu et al., 2016). Similarly, Li et al. (2018); Zha and Li (2019) worked towards a data-less text classification. Veeranna et al. (2016), adopted pre-trained word embedding for measuring semantic similarity between a label and documents. Further endeavor has been spent on zero-shot learning using semantic embeddings by (Hascoet et al., 2019; Zhang et al., 2019; Xie and Virtanen, 2021).

**Uniqueness of our Work:** As mentioned above, the performances of recent LLMs on Bangla documents are mostly unknown. To address this research gap, we examined multiple state-of-the-art sentence encoders and large decoder-based LLMs for a well-defined NLP task, i.e., zero-shot multi-label classification, with an exclusive focus on Bangla documents. Our study lays a foundation for future research endeavors in this under-explored but important direction.

### 3 Problem Statement

We adopt the Definition-Wild 0SHOT-TC methodology proposed by Yin et al. (2019) and later explored by Sarkar et al. (2023a, 2022) for the English language, as defined below.

**Definition 1. 0SHOT-MLC:** *Given a collection of documents denoted as  $D = \{d_1, d_2, \dots, d_n\}$ , a user represented by  $U$  and a set of **user-defined** labels denoted as  $L_U = \{l_1, l_2, \dots, l_m\}$  provided in **real-time**, classify each document  $d_i \in D$  with zero or more labels from  $L_U$ , without further fine-tuning.*

The 0SHOT-MLC employs embeddings of documents and labels, and the detailed methodology for embedding-based zero-shot-MLC is provided in Appendix A.1.

## 4 Experimental Setup

### 4.1 Dataset

We created a new benchmark corpus, BanglaNewsNet, by crawling a large collection of publicly available online news from a renowned Bangladeshi news website, <https://www.prothomalo.com/>. Each article here is already labeled with one or more labels by human annotators. For example, an article titled “মেসিকে রিয়াল বেতিসে চান তাঁর আর্জেন্টিনা দলের স্তীর্থ” (Messi’s Argentina team-mate wants him in Real Betis) is associated with labels “ফুটবল” (Football) and “আর্জেন্টিনা ফুটবল দল” (Argentina Football Team). We scraped the article titles, article texts and labels from the website and further cleaned the dataset by merging similar/repetitive labels into a single label, following approaches in (Sarkar and Karmaker, 2022). See Table 1 for an overview of the dataset.

Dataset	# of Name	Articles	Avg. article length	Labels retained	Labels/ article
BanglaNewsNet	7245	$\approx 2517$ words	21	1.345	

Table 1: An overview of the BanglaNewsNet dataset

### 4.2 Evaluation Metric

For evaluation, we report the Precision, Recall, and  $F_1$  scores. Specifically, for each article, the model-inferred label(s) were compared against the list of “gold” label(s) to compute the true positive, false positive, and false negative statistics, which were then aggregated to compute the final Precision, Recall, and micro-averaged  $F_1$  score.

### 4.3 Classic Encoders Vs. Large Decoders

For baselines, we implemented sentence similarity-based classification using LASER (Artetxe and Schwenk, 2019), LaBSE (Feng et al., 2020), and BanglaTransformer (Muhammad Rafsan Kabir et al., 2024). Additionally, we evaluated large decoder-based models from 13 LLM families, as listed in Table 2. For commercial LLMs (OpenAI, Google, Anthropic), we used their APIs, while open-source models were accessed via the Hugging Face Transformers library. Since APIs for BLOOM (Scao et al., 2022), FLAN-UL2 (Wei et al., 2021) and GPTNeoX (Black et al., 2022) were unavailable, we utilized their embeddings to enable direct comparison with traditional encoders. Beyond embedding-based approaches, we also evaluated LLMs on the Bangla Zero-Shot-MLC task using a prompt-based approach. Unlike embedding similarity methods, which require access to model representations, prompting enables direct multi-label classification using model-generated responses, making it a practical approach when working with black-box LLMs. The complete prompt used for Zero-Shot-MLC task is provided in Table 3.

LLM Family	Model
Google (Team et al., 2024)	gemini-1.5-pro (~200B)
	gemini-1.5-flash (8B)
	gemini-1.0-pro (~200B)
	gemma-2 (27B)
	gemma-2 (9B)
	gemma-2 (2B)
OpenAI (Brown et al., 2020)	gpt-3.5-turbo-0613 (~20B)
	gpt-4o-mini-2024-07-18 (~20B)
	gpt-3 (175B)
EleutherAI (Black et al., 2022)	gpt-neox-20b (20B)
MetaAI Llama3 (Touvron et al., 2023)	Llama-3.1-8B (8B) Llama-3.2-3B (3B) meta-Llama-3.2-11B-Instruct (11B) Llama-3.3-70B-Instruct (70B)
BanglaLlama (Zehady et al., 2024)	BanglaLLama-3.1-8b (8B) BanglaLLama-3.2-3b (3B) BanglaLLama-3-11b (11B)
MistralAI (Jiang et al., 2024)	Mixtral-8x7B-Instruct-v0.1 (56B)
DeepSeek (DeepSeek-AI et al., 2025)	DeepSeek-R1-Distill-Llama-8B (8B) DeepSeek-R1-Distill-Qwen-14B (14B) DeepSeek-R1-Distill-Qwen-32B (32B)
AllenAI (OLMo et al., 2024)	OLMo-2-1124-7B-RM (7B)
Qwen (Yang et al., 2025)	Qwen1.5-72B (72B) Qwen2.5-7B-Instruct-1M (7B)
Gryphe (Gryphe)	MythoMax-L2-13b (13B)
UpStage (Kim et al., 2023)	SOLAR-10.7B-Instruct-v1.0 (10.7B)
Anthropic (Anthropic, 2024)	claude-3-haiku-20240307 (~20B)
BigScience (Scao et al., 2022)	bloom (176B)

Table 2: The list of Large Decoder families and their variants evaluated in Bangla Zero-Shot-MLC.

## 5 Results

This section presents the experimental results on the BanglaNewsNet dataset, evaluating models based on  $F_1$  scores. Table 4 compares embedding-similarity methods using sentence encoders and three large decoder-based models. We utilized two label embedding approaches, such as *Label name + Keywords* and *Explicit-Mentions*, detailed in A.3. Based on the findings of Sarkar et al. (2023b), these embedding methods previously yielded the best results, leading us to focus exclusively on them in this study. Among encoders, LaBSE performed best, achieving an  $F_1$  score of around 40%, while BanglaTransformer performed slightly lower. For large decoder-based models, Flan-UL2, BLOOM, and GPT-NeoX failed to achieve strong results in 0SHOT-MLC on Bangla articles. Although some models exhibited high recall, their low precision made the results less meaningful. Notably, no LLM surpassed LaBSE in embedding-based 0SHOT-MLC, highlighting the limitations of current LLMs in handling their linguistic complexities.

### 5.1 LLM Performance Across Model Sizes

Table 5 presents the performance of LLMs categorized by parameter size. Below is a summary of key observations for each model category:

**Small Models (<8B).** OLMo-7B-Instruct (Groeneweld et al., 2024) and Qwen-2.5-7B-Instruct (Team, 2024b) outperform other models in this category, yet their  $F_1$  scores (0.381 and 0.351, respectively) remain low. Surprisingly, none of the LLMs in this group exceeds the classic encoder LaBSE ( $F_1$ : 0.400), highlighting their limitations in the context of Bangla Zero-Shot-MLC.

**Mid-Sized Models (8B - 10B).** Gemini-1.5-Flash (8B) (Team et al., 2024) ( $F_1$ : 0.571) outperforms larger models, surpassing Gemma-2-9B (Team, 2024a) ( $F_1$ : 0.544) and DeepSeek-R1-Distill-8B (Guo et al., 2025) ( $F_1$ : 0.470).

**Intermediate Models (10B - 15B).** MythoMax-L2 (13B) (Gryphe) underperforms significantly ( $F_1$ : 0.187). Llama 3.2 (11B) (Touvron et al., 2023) and BanglaLlama 3.2 (11B) (Zehady et al., 2024) achieve close scores (0.513 and 0.481), showing minimal benefits from language-specific tuning. MythoMax-L2 (13B) (Gryphe) underperforms significantly ( $F_1$ : 0.187). Llama 3.2 (11B) (Touvron et al., 2023) and BanglaLlama 3.2 (11B) (Zehady et al., 2024) achieve close

Prompt Design	
System setup	
<p>The AI assistant has been designed to understand and categorize user input by the given labels. When processing user input, the assistant must predict the labels from one of the following pre-defined options: 'চাকরিবাকরি' (Job market), 'করোনাভাইরাস' (Coronavirus), 'চলচ্চিত্র ও তারকা' (Movies and celebrities), 'স্বাস্থ্য' (Health), 'ব্যাংক' (Banking), 'অর্থনৈতিক' (Economy), 'শিক্ষা' (Education), 'প্রাকৃতিক দুর্ঘটনা' (Natural disasters), 'আইন ও আদালত' (Law and justice), 'কৃত্ত্বনৈতিক' (Politics), 'শিল্প ও বাণিজ্য' (Industry and commerce), 'ভ্রমণ' (Travel), 'নকশা' (Design), 'ফুটবল' (Football), 'খাবারদাবার' (Food and dining), 'দেশ ও রাজনৈতিক' (Country and politics), 'আন্তর্জাতিক' (International), 'দেশের খবর' (Country news), 'রাশিয়া ইউক্রেন সংঘাত' (Russia-Ukraine conflict), 'ক্রিকেট' (Cricket), 'নারী' (Women). It is essential to note that an article may have multiple labels. If the user input is not relevant with any labels, the assistant should print nothing, indicating that the input does not align with the available categories. The agent MUST response with the following json format: {“Labels”: [“List of labels”]}</p>	
User	<p><b>Taking into account the given Bangla article</b> {ব্যাংক এশিয়া ২০১৪ সালে ব্যাংকিং সেবার বাইরে থাকা বিপুল জনগোষ্ঠীকে ব্যাংকিং সেবায় আনতে এজেন্ট ব্যাংকিং সেবা চালু করে। বর্তমানে এজেন্ট ব্যাংকিং সেবা গ্রহণকারীর সংখ্যা দেড় কোটির বেশি। এর মধ্যে ব্যাংক এশিয়ার গ্রাহক ৫৫ লাখের বেশি। এসব গ্রাহকের ৯২ শতাংশই গ্রামীণ জনগোষ্ঠী। আবার ৬২ শতাংশ গ্রাহকই নারী। সারা দেশে ব্যাংক এশিয়ার ৫ হাজার ৮০০-এর বেশি এজেন্ট আউটলেট রয়েছে, যাদের মধ্যে নারী এজেন্ট ৫৪০ জন।..... (In 2014, Bank Asia introduced agent banking services to bring banking services to a large population that was outside the reach of traditional banking. At present, a total of 31 banks, including both public and private, are providing these services. The number of users availing agent banking services has surpassed 140 million. Among them, Bank Asia has more than 5.5 million customers, of which 92% are from rural areas. Furthermore, 62% of the customers are women. Throughout the country, there are more than 5,400 agent outlets of Bank Asia, including 540 outlets managed by female agents.....)}, predict the category or labels of this article from the list of mentioned labels.</p>
Assistant	{“Labels”: “ব্যাংক” (Banking), “শিল্প ও বাণিজ্য” (Industry and commerce)}
<p><b>Directive:</b> Taking into account the given Bangla article {article text}, predict the category or labels of this article from the list of mentioned {labels}. Please remember to only respond in the predefined JSON format without any additional information.</p>	

Table 3: Prompt design details for the Zero-Shot-MLC task on BanglaNewsNet.

scores (0.513 and 0.481), showing minimal benefits from language-specific tuning. Although BanglaLlama is fine-tuned on Bangla corpora, it fails to outperform the base Llama model. A plausible explanation is that because BanglaLlama was trained on specific data sources such as WikiBangla rather than on task-oriented datasets. As a result, its adaptation improves linguistic alignment but not reasoning ability.

**Larger Models (15B - 50B).** GPT-4o Mini (20B) (Brown et al., 2020) and Claude-3 Haiku (20B) (Anthropic, 2024) perform well ( $F_1$ : 0.588, 0.540). Gemma-2 (27B) (Team, 2024a) leads this category ( $F_1$ : 0.593). DeepSeek R1 Distill (32B) (Guo et al., 2025) also performs well ( $F_1$ : 0.550), reinforcing the effectiveness of model distillation.

**Largest Models (>50B).** Gemini-1.5-Pro (Team et al., 2024) and Gemini-1.0-Pro (Team et al., 2023) achieve the highest  $F_1$  scores (0.616,

0.589). Despite their size, Mixtral-56B-Instruct ( $F_1$ : 0.305) (Jiang et al., 2024), Qwen-1.5 (Yang et al., 2025) ( $F_1$ : 0.429), and GPT-3.5 (Brown et al., 2020) ( $F_1$ : 0.417) significantly underperformed.

Overall, the Gemini models demonstrate strong performance on this task. Gemini-1.5-Pro achieves the highest  $F_1$  score (0.616), driven largely by its high recall (0.918). Although this indicates that the model successfully identifies most relevant instances, such unusually high recall—paired with comparatively low precision (0.463)—suggests that the model may be over-generating labels, capturing correct cases, but also introducing many false positives. In contrast, Gemini-1.0-Pro achieves the second-highest  $F_1$  score (0.589) with notably strong precision (0.796), reflecting more conservative and accurate predictions with fewer false alarms. A full overview

Classical Encoder Models						
	LASER		LaBSE		BanglaTransformer	
Label Embedding ->	Label+KW	Expl.-Ment.	Label+KW	Expl.-Ment.	Label+KW	Expl.-Ment.
F1 Score	0.267	0.305	0.354	0.404	0.334	0.384
Large Decoder Models						
	Flan-UL2		Bloom		GPTNeoX	
Label Embedding ->	Label+KW	Expl.-Ment.	Label+KW	Expl.-Ment.	Label+KW	Expl.-Ment.
F1 Score	0.234	0.241	0.329	0.341	0.345	0.357

Table 4: F1 score comparison of embedding based Zero-Shot-MLC across classical encoders and large decoders.

<8B Parameters		8B - 10B Parameters		10B - 15B Parameters		15B - 50B Parameters		>50B Parameters	
Model	$F_1$	Model	$F_1$	Model	$F_1$	Model	$F_1$	Model	$F_1$
Llama 3.2 (3B)	<b>0.320</b>	Llama 3.1 (8B)	<b>0.476</b>	Llama 3.2 (11B)	<b>0.513</b>	GPT 3.5 Turbo (~20B)	<b>0.470</b>	Llama 3.3 (70B)	<b>0.558</b>
BanglaLlama 3.2 (3B)	<b>0.323</b>	BanglaLlama 3.1 (8B)	<b>0.424</b>	BanglaLlama 3.2 (11B)	<b>0.481</b>	GPT 4o Mini (~20B)	<b>0.588</b>	Gemini 1.0 Pro (~200B)	<b>0.589</b>
Gemma 2 (2B)	<b>0.280</b>	Gemma 1.5 Flash (8B)	<b>0.571</b>	MythoMax L2 (13B)	<b>0.187</b>	Gemma 2 (27B)	<b>0.593</b>	Gemini 1.5 Pro (~200B)	<b>0.616</b>
OLMo 7B Instruct (7B)	<b>0.380</b>	Gemma 2 (9B)	<b>0.544</b>	SOLAR 10.7B Instruct (10.7B)	<b>0.452</b>	Claude 3 Haiku (~20B)	<b>0.540</b>	Mixtral 56B Instruct (56B)	<b>0.305</b>
Qwen 2.5 7B Instruct (7B)	<b>0.351</b>	DeepSeek R1 Distill (8B)	<b>0.470</b>	DeepSeek R1 Distill (14B)	<b>0.495</b>	DeepSeek R1 Distill (32B)	<b>0.550</b>	Qwen 1.5 (72B)	<b>0.429</b>
								GPT 3.5 (175B)	<b>0.537</b>

Table 5: F1 score comparison of prompting-based approaches across models with varying parameter sizes.

of the precision-recall trade-off across all evaluated models is provided in Table 8 in the Appendix.

#### Interesting Findings.

- **Top-performing models.** Gemini-1.5-Pro (200B,  $F_1$ : 0.616) and Gemma-2 (27B,  $F_1$ : 0.593) achieve the highest scores.
- **Bigger isn't always better.** Gemini-1.5-Flash (8B) outperforms larger models like DeepSeek R1 Distill (32B), Mixtral (56B), and GPT-3.5 (175B).
- **Encoders still hold value.** The classic encoder LaBSE ( $F_1$ : .404) surprisingly outperforms all 8B-based models, reinforcing its effectiveness in zero-shot-MLC.
- **LLMs struggle with low-resource languages.** Despite strong performance in English, most models fail to generalize well in Bangla, exposing a multilingual gap.

## 6 Discussion and Conclusion

Our study takes an important step toward evaluating state-of-the-art models for low-resource languages, with a focus on Bangla. We benchmark classic encoders and large decoder-based LLMs for Zero-Shot Multi-Label Classification (Zero-Shot-MLC), uncovering key limitations in both approaches. While large decoders perform well in high-resource languages, their effectiveness in Bangla remains inconsistent, with most failing to achieve reliable classification despite their scale. Classical encoders, though more stable, did not show significant improvement, indicating that neither approach is fully optimized for Bangla and exposing critical gaps in multilingual adaptation.

This work addresses a crucial research gap and also reinforces the need for tailored approaches to handle the complexity of morphologically rich, low-resource languages like Bangla. By evaluating both encoders and decoder-based models, we contribute to improve multilingual NLP for regional languages.

## Limitations

As a limitation, it is worth mentioning that we had restricted access to several large language models (LLMs) such as LaMDA (137B), Jurassic-1 (178B), ERNIE 3.0 Titan (260B), Gopher (280B), GPT-4, Megatron Turing NLG (530B) and DeepSeek-R1 (671B), which limited their inclusion in our experiments. Additionally, while models like ChatGPT were available through an API, their usage was constrained by cost, which scales with model size, making large-scale evaluations financially challenging. Our computational resources were also a limiting factor, as we could only run open-source models up to 70 billion parameters, restricting our ability to test larger models natively. As a result, for models exceeding this threshold, we had to rely solely on API-based access, further increasing cost constraints. Furthermore, hardware limitations affected our ability to

fine-tune models efficiently, limiting our exploration of task-specific optimizations. Another constraint was dataset diversity. Our experiments were conducted on a single dataset, and broader evaluations are necessary to determine whether our findings generalize across different domains.

## 7 Ethical Considerations

In this study we used publicly available news articles to prepare our dataset. We didn't include any personal or sensitive information in our dataset. All the articles we collected are already published online for public reading, and we used them only for research and analysis. We evaluated different language models which may carry hidden biases. We did not fine-tune any commercial model with private data, and all API-based evaluations followed the provider's usage policies. As our experiment focused on Bangla, a low-resource language, models may show uneven performance across different topics and lead to misclassifications. We encouraged responsible use of our findings and advise that any downstream use should consider these limitations.

## References

Md Abdul Mazid Adnan and Md Musfiq Anwar. 2022. Progressive guidance categorization using transformer-based deep neural network architecture. In *Hybrid Intelligent Systems: 21st International Conference on Hybrid Intelligent Systems (HIS 2021), December 14–16, 2021*, volume 420, page 344. Springer Nature.

Md. Habibul Alam, Md-Mizanur Rahoman, and Md. Abul Kalam Azad. 2017. *Sentiment analysis for bangla sentences using convolutional neural network*. In *2017 20th International Conference of Computer and Information Technology (ICCIT)*, pages 1–6.

Tanvirul Alam, Akib Khan, and Firoj Alam. 2020. Bangla text classification using transformers. *arXiv preprint arXiv:2011.04446*.

Anthropic. 2024. *The claude 3 model family: Opus, sonnet, haiku*. Accessed: 2025-02-14.

Md Musfiq Anwar, Mohammad Zabed Anwar, and M Al-Amin Bhuiyan. 2009. Syntax analysis and machine translation of bangla sentences. *International Journal of Computer Science and Network Security*, 9(8):317–326.

Mikel Artetxe and Holger Schwenk. 2019. *Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond*. *Trans. Assoc. Comput. Linguistics*, 7:597–610.

Tanjim Taharat Aurpa, Richita Khandakar Rifat, Md Shoaib Ahmed, Md Musfiq Anwar, and ABM Shawkat Ali. 2022a. Reading comprehension based question answering system in bangla language with transformer-based learning. *Helicon*, 8(10):e11052.

Tanjim Taharat Aurpa, Rifat Sadik, and Md Shoaib Ahmed. 2022b. Abusive bangla comments detection on facebook using transformer-based deep learning models. *Social Network Analysis and Mining*, 12(1):24.

Nayan Banik and Md. Hasan Hafizur Rahman. 2018. *Gru based named entity recognition system for bangla online newspapers*. In *2018 International Conference on Innovation in Engineering and Technology (ICIET)*, pages 1–6.

Abhik Bhattacharjee, Tahmid Hasan, Wasi Ahmad, Kazi Samin Mubasshir, Md Saiful Islam, Anindya Iqbal, M. Sohel Rahman, and Rifat Shahriyar. 2022. *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 Shahri-

yar. 2023. BanglaNLG and BanglaT5: Benchmarks and resources for evaluating low-resource natural language generation in Bangla. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 726–735, Dubrovnik, Croatia. Association for Computational Linguistics.

Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Lawrence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Bidyut Baran Chaudhuri and Suvankar Bhattacharya. 2008. An experiment on automatic detection of named entities in Bangla. In *Proceedings of the IJCNLP-08 Workshop on Named Entity Recognition for South and South East Asian Languages*.

Amit Das, Md Najib Hasan, Souvika Sarkar, Zheng Zhang, Fatemeh Jamshidi, Tathagata Bhattacharya, Nilanjana Raychowdhury, Dongji Feng, Vinija Jain, and Aman Chadha. 2025. Investigating hallucination in conversations for low resource languages. *arXiv preprint arXiv:2507.22720*.

Amit Das, Zheng Zhang, Najib Hasan, Souvika Sarkar, Fatemeh Jamshidi, Tathagata Bhattacharya, Mostafa Rahgouy, Nilanjana Raychowdhary, Dongji Feng, Vinija Jain, et al. 2024. Investigating annotator bias in large language models for hate speech detection. *arXiv preprint arXiv:2406.11109*.

Daya Guo DeepSeek-AI, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Argha Chandra Dhar, Arna Roy, Md Ahsan Habib, MAH Akhand, and N Siddique. 2022. Transformer deep learning model for bangla–english machine translation. In *Proceedings of 2nd International Conference on Artificial Intelligence: Advances and Applications: ICAIAA 2021*, pages 255–265. Springer.

Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. Language-agnostic bert sentence embedding. *arXiv preprint arXiv:2007.01852*.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.

Gryphe. Mythomax-l2-13b-turbo. <https://deepinfra.com/Gryphe/MythoMax-L2-13b-turbo/api>. Accessed: 2025-02-14.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Md. Arid Hasan, Firoj Alam, Shammur Absar Chowdhury, and Naira Khan. 2019. Neural machine translation for the bangla-english language pair. In *2019 22nd International Conference on Computer and Information Technology (ICCIT)*, pages 1–6.

Tristan Hascoet, Yasuo Ariki, and Tetsuya Takiguchi. 2019. Semantic embeddings of generic objects for zero-shot learning.

*EURASIP Journal on Image and Video Processing*, 2019(1):1–14.

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.

Md Zobaer Hossain, Md Ashraful Rahman, Md Saiful Islam, and Sudipta Kar. 2020. **BanFakeNews: A dataset for detecting fake news in Bangla**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2862–2871, Marseille, France. European Language Resources Association.

Md Gulzar Hussain, Md Rashidul Hasan, Mahmuda Rahman, Joy Protim, and Sakib Al Hasan. 2020. **Detection of bangla fake news using mnb and svm classifier**. In *2020 International Conference on Computing, Electronics & Communications Engineering (iCCECE)*, pages 81–85.

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, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Md. Aminul Islam, Md. Fasihul Kabir, Khan-daker Abdullah-Al-Mamun, and Mohammad Nurul Huda. 2016. **Word/phrase based answer type classification for bengali question answering system**. In *2016 5th International Conference on Informatics, Electronics and Vision (ICIEV)*, pages 445–448.

Tanvirul Islam, Subhenur Latif, and Nadim Ahmed. 2019. **Using social networks to detect malicious bangla text content**. In *2019 1st International Conference on Advances in Science, Engineering and Robotics Technology (ICASERT)*, pages 1–4.

Sabir Ismail, M. Shahidur Rahman, and Md Abdullah Al Mumin. 2014. Developing an automated bangla parts of speech tagged dictionary. *16th Int'l Conf. Computer and Information Technology*, pages 355–359.

Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. **Mixtral of experts**. *arXiv preprint arXiv:2401.04088*.

Wenxiang Jiao, Wenzuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. **Is chatgpt a good translator? a preliminary study**. *arXiv preprint arXiv:2301.08745*.

Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. **The state and fate of linguistic diversity and inclusion in the NLP world**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.

Shubhra Kanti Karmaker Santu, Parikshit Sondhi, and ChengXiang Zhai. 2016. Generative feature language models for mining implicit features from customer reviews. In *Proceedings of the 25th ACM international conference on information and knowledge management*, pages 929–938.

Dahyun Kim, Chanjun Park, Sanghoon Kim, Wonsung Lee, Wonho Song, Yunsu Kim, Hyeyonwoo Kim, Yungi Kim, Hyeyonju Lee, Jihoo Kim, et al. 2023. **Solar 10.7 b: Scaling large language models with simple yet effective depth up-scaling**. *arXiv preprint arXiv:2312.15166*.

M. Kowsher, Abdullah As Sami, Nusrat Jahan Prottasha, Mohammad Shamsul Arefin, Pranab Kumar Dhar, and Takeshi Koshiba. 2022. **Bangla-bert: Transformer-based efficient model for transfer learning and language understanding**. *IEEE Access*, 10:91855–91870.

Md. Kowsher, M M Mahabubur Rahman, Sk Shohorab Ahmed, and Nusrat Jahan Protasha. 2019. **Bangla intelligence question answering system based on mathematics and statistics**. In *2019 22nd International Conference on Computer and Information Technology (ICCIT)*, pages 1–6.

Ximing Li, Changchun Li, Jinjin Chi, Jihong Ouyang, and Chenliang Li. 2018. **Dataless text classification: A topic modeling approach with document manifold**. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, page 973–982, New York, NY, USA. Association for Computing Machinery.

Muhammad Rafsan Kabir, Md. Mohibur Rahman Nabil, and Mohammad Ashrafuzzaman Khan. 2024. **Banglaembed: Efficient sentence embedding models for a low-resource language using cross-lingual distillation techniques**. *arXiv preprint arXiv:2411.15270*. Preprint.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, et al. 2024. **2 olmo 2 furious**. *arXiv preprint arXiv:2501.00656*.

Md Aminul Haque Palash, Md Abdullah Al Nasim, Sourav Saha, Faria Afrin, Raisa Mallik, and Sathishkumar Samiappan. 2022. **Bangla image caption generation through cnn-transformer based encoder-decoder network**. In *Proceedings of International Conference on Fourth Industrial Revolution and Beyond 2021*, pages 631–644. Springer.

Mohammad Osiur Rahman, Fouzia Asharf Mousumi, Edgar Scavino, Aini Hussain, and Hassan Basri. 2008. **Real time road sign recognition system using artificial neural networks for bengali textual information box**. In *2008 International Symposium on Information Technology*, volume 2, pages 1–8.

Nils Reimers and Iryna Gurevych. 2019. **Sentence-bert: Sentence embeddings us-**ing siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3973–3983.

Malik Sallam. 2023. **Chatgpt utility in health-care education, research, and practice: systematic review on the promising perspectives and valid concerns**. In *Healthcare*, volume 11, page 887. MDPI.

Souviaka Sarkar, Dongji Feng, and Shubhra Kanti Karmaker Santu. 2022. **Exploring universal sentence encoders for zero-shot text classification**. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 135–147, Online only. Association for Computational Linguistics.

Souviaka Sarkar, Dongji Feng, and Shubhra Kanti Karmaker Santu. 2023a. **Zero-shot multi-label topic inference with sentence encoders and LLMs**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16218–16233, Singapore. Association for Computational Linguistics.

Souviaka Sarkar, Dongji Feng, and Shubhra Kanti Karmaker\_Santu. 2023b. **Zero-shot multi-label topic inference with sentence encoders and llms**. Association for Computational Linguistics.

Souviaka Sarkar and Shubhra Kanti Karmaker. 2022. **Concept annotation from users perspective: A new challenge**. In *Companion proceedings of the web conference 2022*, pages 1180–1188.

Sourav Sarker, Syeda Tamanna Alam Monisha, and Md Mahadi Hasan Nahid. 2019. **Bengali question answering system for factoid questions: A statistical approach**. In *2019 International Conference on Bangla*

*Speech and Language Processing (ICBSLP)*, pages 1–5.

Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.

S M A Sharif, Nabeel Mohammed, Nafees Mansoor, and Sifat Momen. 2016. *A hybrid deep model with hog features for bangla handwritten numeral classification*. In *2016 9th International Conference on Electrical and Computer Engineering (ICECE)*, pages 463–466.

Nigar M Shafiq Surameery and Mohammed Y Shakor. 2023. Use chat gpt to solve programming bugs. *International Journal of Information Technology & Computer Engineering (IJITC) ISSN: 2455-5290*, 3(01):17–22.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

Gemma Team. 2024a. *Gemma*.

Qwen Team. 2024b. Qwen2. 5: A party of foundation models. *Qwen (Sept. 2024)*. url: <https://qwenlm.github.io/blog/qwen2.5>.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Abdul Hasib Uddin, Sumit Kumar Dam, and Abu Shamim Mohammad Arif. 2019. Extracting severe negative sentence pattern from bangla data via long short-term memory neural network. In *2019 4th International Conference on Electrical Information and Communication Technology (EICT)*, pages 1–6.

Sappadla Prateek Veeranna, Jinseok Nam, EL Mencía, and J Furnkranz. 2016. Using semantic similarity for multi-label zero-shot classification of text documents. In *Proceeding of European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges, Belgium: Elsevier*, pages 423–428.

Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

Huang Xie and Tuomas Virtanen. 2021. Zero-shot audio classification via semantic embeddings. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1233–1242.

An Yang, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoyan Huang, Jiandong Jiang, Jianhong Tu, Jianwei Zhang, Jingren Zhou, et al. 2025. Qwen2. 5-1m technical report. *arXiv preprint arXiv:2501.15383*.

Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3912–3921. Association for Computational Linguistics.

Abdullah Khan Zehady, Safi Al Mamun, Naymul Islam, and Santu Karmaker. 2024. Bongllama: Llama for bangla language. *arXiv preprint arXiv:2410.21200*.

Daochen Zha and Chenliang Li. 2019. Multi-label dataless text classification with topic modeling. *Knowledge and Information Systems*, 61(1):137–160.

Jingqing Zhang, Piyawat Lertvittayakumjorn, and Yike Guo. 2019. [Integrating semantic knowledge to tackle zero-shot text classification](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 27, 2019, Volume 1 (Long and Short Papers)*, pages 1031–1040. Association for Computational Linguistics.

## A Appendix

### A.1 Encoder & LLM Based 0-shotMLC

Here, we describe the steps used in our 0-shot-MLC approach.

1. Input Document: The end user provides the article text, custom-defined labels, and a set of keywords (optional).
2. Embedding Generation: The article text, labels, and keywords are transformed into rich embedding by leveraging the power of language models and encoding methods, capturing the essence of their content.
  - **Article Embedding**: We embed the entire article with sentence encoders and LLMs in a single shot.
  - **Label Embedding**: We adopted two different ways for target label embedding: 1) *Label + Keyword*- Label embedding using label name and keywords, 2) *Explicit-Mentions*- Label embedding using article-text which contains explicit mentions of label names.

The details of these embeddings have been discussed in the Appendix A.2, A.3. While there are certainly other embedding methods possible, based on the findings of the paper (Sarkar et al., 2023b), these embedding combinations worked best previously,

and hence, we focused only on the above embedding type.

3. Threshold-based Label Assignment: Next, we quantify the cosine similarity between the article embedding and the label embeddings. Labels are assigned to the article based on a specified threshold, indicating the presence of a strong association. By experimenting with different threshold values (ranging from 0-1), a comprehensive analysis is conducted.
4. Zero-Shot multi-label classification: The outcome of this classifier is the prediction of relevant label(s) for the given article.

### A.2 “Entire Article” based embedding

Encode the entire article using sentence encoders or LLMs in a single shot, including articles that are long paragraphs and consist of more than one sentence.

### A.3 Label Embedding Approaches

We have used 2 different approaches for computing label embedding. The consecutive sections discuss about different procedures for generating label embedding.

#### A.3.1 “Label name + Keywords” based embedding

Encode both label name and keywords, then average all embeddings to generate the final label embedding.

#### A.3.2 “Explicit-Mentions” based embedding

First, we extract all the articles explicitly mentioning the label/phrase using algorithm 1 for all labels. Then, for each label, we generate embeddings of all articles that are explicitly annotated/classified with that label, then average them to obtain the ultimate label embedding.

### A.4 Baseline Sentence Encoders

This section presents a bird’s-eye view of the sentence encoders and large language models we have used for our experiments.

- **Language-Agnostic SEntence Representations (LASER)**: LASER is

---

**Algorithm 1** Article Annotation using Explicit Mention

---

```
1: Input: Article text, Label names, and Keywords
2: Output: Articles annotated with explicit Label
3: for each article text do
4:   check whether the label name or set (at least 3) of
      the informative keywords are present or not in the
      corresponding article text
5:   if present then annotate the article with the explicit
      label
6:   end if
7: end for
```

---

a sentence encoding model that generates language-agnostic representations (Artetxe and Schwenk, 2019). It is capable of encoding sentences from multiple languages into fixed-length vectors, enabling cross-lingual tasks and multilingual applications.

- **Language-agnostic BERT Sentence Embedding (LaBSE):** LaBSE (Feng et al., 2020) is a language-agnostic model based on the BERT architecture. It provides sentence embeddings that capture the semantic meaning of sentences across different languages. LaBSE allows for cross-lingual understanding and transfer-learning tasks.
- **Bangla sentence embedding transformer:** This Bangla sentence transformer (Muhammad Rafsan Kabir et al., 2024) is specifically designed for the Bangla language. It utilizes a transformer-based architecture to encode Bangla sentences into meaningful representations, enabling various NLP tasks in Bangla text analysis. It was trained on 2,50,000 Bangla sentences(wiki) by Sentence-Transformer. This work is inspired by Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks (Reimers and Gurevych, 2019) technique.

## A.5 Large Language Models

- **BLOOM:** Scao et al. (2022) introduce BLOOM, a massive language model with 176 billion parameters. BLOOM is trained on 46 natural languages and 13 programming languages and is the result of a collaborative effort involving

hundreds of researchers. BLOOM is a causal language model trained to predict the next token in a sentence. This approach has been found effective in capturing reasoning abilities in large language models. BLOOM uses a Transformer architecture composed of an input embedding layer, 70 Transformer blocks, and an output language-modeling layer. The sequential operation of predicting the next token involves passing the input tokens through each of the 70 BLOOM blocks. To prevent memory overflow, only one block is loaded into RAM at a time. The word embeddings and output language-modeling layer can be loaded on demand from disk.

- **GPT-NeoX:** The GPT-NeoX-20B paper, authored by the Black et al. (2022), introduces an architecture similar to GPT-3 but with notable differences. They utilize rotary positional embeddings for token position encoding instead of learned embeddings and parallelize the attention and feed-forward layers, resulting in a 15% increase in throughput. Unlike GPT-3, GPT-NeoX-20B exclusively employs dense layers. The authors trained GPT-NeoX-20B using EleutherAI’s custom codebase (GPT-NeoX) based on Megatron and DeepSpeed, implemented in PyTorch. To address computational limitations, the authors reused the hyperparameters from the GPT-3 paper. In their evaluation, the researchers compared GPT-NeoX-20B’s performance to their previous model, GPT-J-6B, as well as Meta’s FairSeq 13B and different sizes of GPT-3 on various NLP benchmarks, including LAMBADA, Winogrande, HendrycksTest, and the MATH dataset. While improvements were desired for NLP tasks, GPT-NeoX-20B exhibited exceptional performance in science and math tasks.
- **Google (Gemini & Flan):** (Team et al., 2024) introduced Gemini-1.5 Pro and 1.0 Pro models containing powerful multimodal capabilities, allowing them to process text, images, audio, and video at the

same time. This makes them highly effective for real-time AI tasks like reasoning and code generation. With an optimized transformer-decoder and sparse mixture-of-experts techniques, they enhance efficiency while handling complex inputs. Their long-context attention feature also helps them retain and process longer pieces of information more effectively. Meanwhile, the Gemma family (2B, 9B, and 27B) is built for flexible deployment, whether on-device or in the cloud, making it ideal for fast, low-latency applications. Flan-UL2, an improved version of Flan-T5, uses Mixture-of-Denoisers (MoD) pre-training to push the boundaries of NLP tasks like classification, reasoning, and question answering. With 20 billion parameters, it outperforms models like T5 and GPT-3, excelling in zero-shot learning and chain-of-thought reasoning while achieving top results on major NLP benchmarks.

- **OpenAI (GPT Series):** According to (Brown et al., 2020) GPT-4o, GPT-3.5 Turbo, and GPT-3 represent major leaps in AI language models. GPT-4o is the most advanced, handling text, images, and audio at the same time with better speed and accuracy, making it great for real-time AI applications like chat-bots and coding assistants. GPT-3.5 Turbo is designed for efficiency, balancing cost and performance, which makes it popular for business AI tools and content generation. GPT-3, with its 175 billion parameters, was a game-changer in AI, setting the stage for today's models with its strong language understanding and reinforcement learning for better alignment with human values.
- **MetaAI (Llama3):** (Touvron et al., 2023) mentioned Meta's Llama3 models Llama-3.1-8B, Llama-3.2-3B, and Llama-3.3-70B are designed for efficient, cost-effective AI deployment. They use adaptive tokenization and transformer pruning to reduce computational demands while maintaining strong performance. The Llama-3.2-11B Vision-

Instruct model expands Meta's work in multi-modal AI, integrating visual and language reasoning for applications like computer vision, medical imaging, and smart assistants. By open-sourcing its models, Meta promotes collaborative AI development and innovation in decentralized AI systems. With real-time processing capabilities, Llama3 models stand out for their speed and energy efficiency, making them ideal for on-premise AI, embedded systems, and low-latency conversational tools.

- **BanglaLlama:** (Zehady et al., 2024) introduced BanglaLlama which is a groundbreaking initiative aimed at improving NLP for the Bangla language, addressing the lack of high-quality AI models for low-resource languages. With models ranging from 3B to 11B parameters, it is trained specifically on Bangla text to better capture linguistic nuances, cultural context, and accuracy. By using advanced tokenization and dataset augmentation, BanglaLlama excels in translation, content creation, and conversational AI. Its fine-tuned instruction-following capabilities help make AI more inclusive, ensuring non-English languages are better represented in global AI advancements. Beyond language processing, BanglaLlama also plays a crucial role in reducing biases, adapting to different Bangla dialects, and preserving indigenous languages in AI systems.
- **MistralAI:** (Jiang et al., 2024) represents Mixtral-8x7B as a groundbreaking Mixture of Experts (MoE) model, significantly reducing computational complexity while enhancing inference efficiency. Unlike conventional transformer architectures, Mixtral activates only a subset of its parameters per forward pass, reducing memory footprint and improving scalability. This innovative architecture allows highly efficient parallelized computation, making it a top contender for large-scale enterprise AI applications, real-time NLP solutions, and multilingual text generation. Mis-

tralAI’s MoE-based LLMs are widely recognized for their superior speed-to-accuracy trade-offs, positioning them as one of the most energy-efficient large-scale models in the industry.

- **DeepSeek:** (DeepSeek-AI et al., 2025) mentioned DeepSeek-R1 series, is designed for maximum efficiency, using advanced knowledge distillation to maintain strong performance with fewer parameters. With techniques like quantization and structured pruning, these models deliver fast inference speeds and low-latency processing, making them perfect for AI applications with limited resources. Ideal for real-time NLP, enterprise automation, and AI assistants, DeepSeek-R1 ensures quick responses while keeping computational demands low, making it a great choice for businesses and developers focused on efficiency.
- **AllenAI:** (OLMo et al., 2024) designed OLMo-2-7B with a strong focus on explainable AI (XAI) and interpretability. Unlike traditional black-box AI models, it incorporates features like attention transparency and explainability layers, ensuring clearer insights into how it processes information. Optimized for research in linguistics, AI ethics, and decision-making, OLMo is a top choice for academics, policymakers, and those building transparent AI systems that prioritize human understanding and trust.
- **Alibaba:** According to (Yang et al., 2025) Alibaba’s Qwen series, including Qwen1.5-72B and Qwen2.5-7B, is built for enterprise applications, specializing in complex reasoning, structured NLP, and industry-specific adaptability. These models are fine-tuned for tasks like financial analysis, healthcare AI, and multilingual document processing, making them highly versatile for business use. With strong instruction tuning, the Qwen models excel at knowledge-intensive tasks, delivering high factual accuracy and outperforming many mainstream models in industry automation and specialized AI

applications.

- **Gryphe:** (Gryphe) designed Gryphe’s MythoMax-L2-13B for creative content generation, interactive storytelling, and narrative coherence. This model incorporates fine-tuned stylistic awareness and logical consistency, making it a preferred choice for conversational AI, virtual assistants, and AI-driven journalism.
- **UpStage:** (Kim et al., 2023) designed SOLAR-10.7B for flexibility, excelling in instruction tuning, few-shot learning, and adapting to various AI tasks. With a structured fine-tuning approach, it performs exceptionally well in areas like legal text interpretation, scientific research support, and summarizing complex documents with context and accuracy.
- **Anthropic:** (Anthropic, 2024) built Claude-3 Haiku with a strong focus on ethical AI, safety, and human-guided learning (RLHF). It excels at maintaining context, ensuring fairness, and delivering reliable real-time conversations, making it a great fit for critical applications in healthcare, governance, and compliance-focused AI.

## A.6 Precision-Recall Trade-Off in LLMs for Bangla

The performance analysis of large language models (LLMs) for zero-shot multi-label classification (MLC) in Bangla highlights several important aspects, particularly the trade-offs between precision and recall, which we have explored in the appendix due to space constraints.

### A.6.1 Sentence-Encoder Models

Looking at the baseline sentence-encoder-based approaches (Table 6), we see that BanglaTransformer achieves the highest  $F_1$  score (0.334), but with an evident imbalance between precision and recall. This pattern is also observed in LASER and LABSE, where recall is consistently higher than precision. While higher recall means the model retrieves more relevant labels, it also increases the number of false positives, which is a common issue in low-resource languages where high-quality training data is limited.

Topic+Keywords Based Label Embedding									
LASER			LaBSE			BanglaTransformer			
Precision	Recall	$F_1$	Precision	Recall	$F_1$	Precision	Recall	$F_1$	
0.162	0.750	0.267	0.282	0.477	<b>0.354</b>	0.224	0.648	0.334	
Explicit-Mention Based Label Embedding									
LASER			LaBSE			BanglaTransformer			
Precision	Recall	$F_1$	Precision	Recall	$F_1$	Precision	Recall	$F_1$	
0.193	0.724	0.305	0.300	0.617	<b>0.404</b>	0.276	0.635	0.384	

Table 6: Performance comparison of baseline sentence encoder-based approaches.

Topic+Keywords Based Label Embedding									
FLAN-UL2			BLOOM			GPT-NeoX			
Precision	Recall	$F_1$	Precision	Recall	$F_1$	Precision	Recall	$F_1$	
0.135	0.890	0.234	0.231	0.574	0.329	0.235	0.634	0.345	
Explicit-Mention Based Label Embedding									
FLAN-UL2			BLOOM			GPT-NeoX			
Precision	Recall	$F_1$	Precision	Recall	$F_1$	Precision	Recall	$F_1$	
0.144	0.742	0.241	0.232	0.642	0.341	0.241	0.675	0.357	

Table 7: Performance comparison of different large language models.

Performance Comparison of LLMs with Varying Parameter Sizes						
Parameter Size	Model	Precision	Recall	$F_1$		
<8B	Llama 3.2 (3B)	0.203	0.710	0.320		
	BanglaLlama 3.2 (3B)	0.219	0.595	0.323		
	Gemma 2 (2B)	0.164	0.942	0.280		
	OLMo 7B Instruct (7B)	0.365	0.410	0.380		
	Qwens 2.5 7B Instruct (7B)	0.218	0.906	0.351		
8B to 10B	Llama 3.1 (8B)	0.325	0.886	0.476		
	BanglaLlama 3.1 (8B)	0.289	0.790	0.424		
	Gemini 1.5 Flash (8B)	0.416	0.905	0.571		
	Gemma 2 (9B)	0.388	0.913	0.544		
	DeepSeek R1 Distill (8B)	0.343	0.782	0.470		
10B to 15B	Llama 3.2 (11B)	0.359	0.895	0.513		
	BanglaLlama 3.2 (11B)	0.336	0.849	0.481		
	MythoMax L2 (13B)	0.128	0.347	0.187		
	SOLAR 10.7B Instruct (10.7B)	0.342	0.650	0.452		
	DeepSeek R1 Distill (14B)	0.350	0.850	0.495		
15B to 50B	GPT 3.5 Turbo (~20)	0.350	0.741	0.470		
	GPT 4o Mini (~20B)	0.439	0.889	0.588		
	Gemma 2 (27B)	0.450	0.833	0.593		
	Claude 3 Haiku (~20B)	0.402	0.810	0.540		
	DeepSeek R1 Distill (32B)	0.418	0.804	0.550		
50B>	Llama 3.3 (70B)	0.401	0.921	0.558		
	Gemini 1.0 Pro (~200B)	0.796	0.468	0.589		
	Gemini 1.5 Pro (~200B)	0.463	0.918	0.616		
	Mixtral 56B Instruct (56B)	0.194	0.630	0.305		
	Qwens 1.5 (72B)	0.353	0.540	0.429		
	GPT 3.5 (175B)	0.515	0.573	0.537		

Table 8: Performance comparison of prompting-based approaches across models with varying parameter sizes.

### A.6.2 Large-Scale Generative Models

Shifting to large-scale generative models (Table 7), we observe a noticeable improvement in  $F_1$  scores, with GPT-NeoX (Black et al., 2022) ( $F_1$ : 0.357) outperforming BLOOM and FLAN-UL2 (Wei et al., 2021). However, GPT-NeoX’s (Black et al., 2022) precision (0.241) remains significantly lower than its recall (0.675), reinforcing a major trend seen in generative models: they tend to be highly recall-biased, favoring coverage over accuracy. This can be attributed to their training approach, which optimizes for broad knowledge retrieval rather than precise classification. While this makes them effective for open-ended generation tasks, it poses a challenge for multi-label classification, where specificity is crucial. A recall-heavy approach may work well in some cases, such as medical document classification, where missing a critical label could be costly, but for general-purpose classification tasks, such a model could introduce significant noise. This reinforces the need for additional fine-tuning or hybrid methods to improve precision without sacrificing recall. Examining LLMs across different parameter sizes (Table 8), we see a clear scaling trend, where larger models generally achieve higher  $F_1$  scores but with increasing recall at the expense of precision. Gemini 1.5 Pro (200B) (Team et al., 2024) achieves the highest  $F_1$  score (0.616), with a recall of 0.918 and precision of just 0.463. This means that while it effectively captures relevant labels, it also introduces substantial noise in classification. A similar pattern is seen with GPT-4o Mini (20B) (Brown et al., 2020) and Claude-3 Haiku (20B) (Anthropic, 2024), with recall values of 0.889 and 0.810, respectively. These models demonstrate strong generalization and retrieval capabilities but lack the specificity required for accurate multi-label classification. However, Gemma 2 (27B) (Team et al., 2024) and DeepSeek R1 Distill (32B) (DeepSeek-AI et al., 2025) achieve more balanced trade-offs, with precision scores of 0.450 and 0.418 and recall scores of 0.833 and 0.804, respectively. This suggests that well-optimized architectures and distillation techniques can enable mid-sized

models to match or even surpass larger models in classification tasks. Interestingly, GPT-3.5 (175B) (Brown et al., 2020) underperforms with an  $F_1$  score of 0.537, reinforcing that parameter size alone does not guarantee superior classification accuracy. This aligns with previous findings in NLP research, where training data quality, fine-tuning strategies, and task-specific optimizations often play a more significant role than raw model size.

### A.6.3 Final Thoughts and Key Findings

From a statistical standpoint, the precision-recall trade-off highly indicative of models underlying architectures and training methodologies. Sentence encoders, effective in recall-driven tasks, but fail to deliver precise classifications due to their limited exposure to label dependencies. Instruction-tuned LLMs, on the other hand, benefit from broader generalization but often lack the necessary specificity for multi-label classification, leading to recall-heavy biases. Notably, distilled models like DeepSeek R1 Distill (32B) (DeepSeek-AI et al., 2025) demonstrate a more balanced performance, suggesting that parameter-efficient architectures can even outperform, larger models.