

# BURMESE-SAN: Burmese NLP Benchmark for Evaluating Large Language Models

Thura Aung<sup>1,\*</sup>, Jann Railey Montalan<sup>2,3</sup>, Jian Gang Ngui<sup>2,3</sup>, Peerat Limkonchotiwat<sup>2,3</sup>

<sup>1</sup>King Mongkut's Institute of Technology Ladkrabang

<sup>2</sup>AI Singapore <sup>3</sup> National University of Singapore

66011606@kmitl.ac.th

{railey, jiangangngui, peerat.limkonchotiwat}@aisingapore.org

## Abstract

We introduce *BURMESE-SAN*, the first holistic benchmark that systematically evaluates large language models (LLMs) for Burmese across three core NLP competencies: understanding (NLU), reasoning (NLR), and generation (NLG). *BURMESE-SAN* consolidates seven subtasks spanning these competencies, including Question Answering, Sentiment Analysis, Toxicity Detection, Causal Reasoning, Natural Language Inference, Abstractive Summarization, and Machine Translation, several of which were previously unavailable for Burmese. The benchmark is constructed through a rigorous native-speaker-driven process to ensure linguistic naturalness, fluency, and cultural authenticity while minimizing translation-induced artifacts. We conduct a large-scale evaluation of both open-weight and commercial LLMs to examine challenges in Burmese modeling arising from limited pretraining coverage, rich morphology, and syntactic variation. Our results show that Burmese performance depends more on architectural design, language representation, and instruction tuning than on model scale alone. In particular, southeast asia regional fine-tuning and newer model generations yield substantial gains. Finally, we release *BURMESE-SAN* as a public leaderboard to support systematic evaluation and sustained progress in Burmese and other low-resource languages. <https://leaderboard.sea-lion.ai/detailed/MY>.

**Keywords:** Burmese NLP, Large Language Models, Benchmark, Low-Resource Language

## 1. Introduction

Recent advances in Large Language Models (LLM) development have improved the overall capabilities of Natural Language Processing (NLP). With a large amount of data and high computing power, LLMs have seen widespread adoption due to their ability to provide scalable, practical, and diverse capabilities (Guo et al., 2023; Hadi et al., 2023). In particular, studies (Liang et al., 2023; Qin et al., 2025; Chang et al., 2024) demonstrate the effectiveness of LLMs across various languages and tasks, utilizing their benchmarks to evaluate and compare model performance.

HELM (Liang et al., 2023) proposed formulating a holistic benchmark to evaluate the robustness of LLMs for a diversity of tasks in English. SEACrowd (Lovenia et al., 2024) and SEA-HELM (Susanto et al., 2025), designed as holistic benchmarks, evaluate LLMs' linguistic capabilities in several Southeast Asian (SEA) languages, and both use human-crafted and machine-generated corpora. The fact that some LLMs can perform well on these benchmarks demonstrates that LLMs can generalize from English to SEA languages.

However, most existing LLM benchmarks focus on English, with relatively few evaluating non-English languages (Son et al., 2024). Furthermore, even in multilingual benchmarks (Liang et al., 2023;

Susanto et al., 2025), Burmese is not included because the necessary datasets are not yet available. Despite being the official language of Myanmar and having a total number of speakers of over 43-45 million (Ethnologue, 2019), Burmese remains an under-resourced language for both models and benchmarks (Dou et al., 2025).

Burmese stands out among the major SEA languages for its extensive system of case marking and its rich morphological structure (Jenny, 2021). It is a tonal language with an analytic grammatical structure (Okell, 2023). Burmese features four distinct tones, with each syllable carrying a tone that plays a crucial role in distinguishing meaning between otherwise identical syllables.

The language follows a subject-object-verb (SOV) word order and utilizes a complex system of honorifics and politeness markers that reflect social hierarchy and relationships (Republic of the Union of Myanmar, Myanmar Language Commission, 2013; Khin Aye, 2024). Burmese is also a diglossic language, with formal and informal varieties that can differ significantly in vocabulary and structure, adding another layer of complexity for language modeling and evaluation. Therefore, due to the challenge of language and the availability of Burmese corpora, we need to dedicate more resources to creating a high-quality benchmark to identify any gaps and challenges in current LLMs.

To tackle these challenges, we propose

\* Work conducted during Research Internship at AI Singapore, National University of Singapore.

# BURMESE-SAN

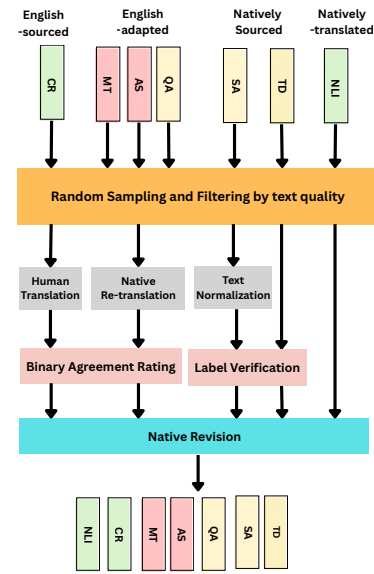
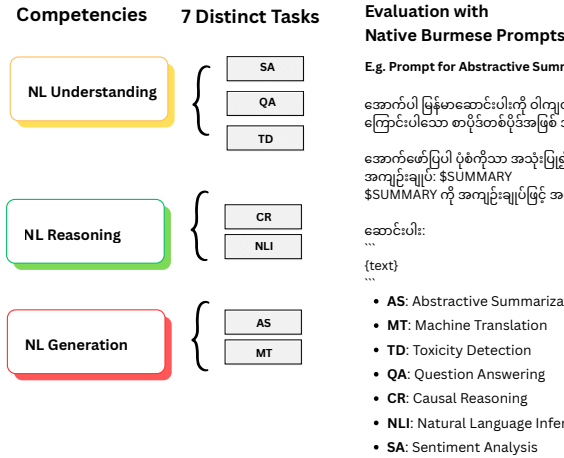


Figure 1: **BURMESE-SAN Benchmark (Left) and Dataset Curation Process for the benchmark (Right)**. *BURMESE-SAN* is a benchmark that holistically evaluates LLM performance across a wide range of Burmese language tasks. The evaluation is based on native Burmese text, with prompts written in formal Burmese to ensure clarity and grammatical correctness.

*BURMESE-SAN*<sup>1</sup>, the first holistic Burmese Benchmark for evaluating LLMs across seven NLP tasks covering natural language understanding, reasoning, and generation. In particular, as shown in Figure 1, we have seven subtasks: sentiment analysis (SA), toxicity detection (TD), question answering (QA), causal reasoning (CR), natural language inference (NLI), abstractive summarization (AS), and machine translation (MT), with a total of 3,920 samples. *BURMESE-SAN* was meticulously built in collaboration with native speakers, who have been involved at every stage of the pipeline - from task selection, prompt translation, and data validation to final quality assurance. This high-effort, *human-centered* approach ensures linguistic authenticity, cultural relevance, and evaluative rigor, setting a new standard for Burmese benchmarks.

**Proposed Studies.** Based on *BURMESE-SAN*, we investigate the following research questions:

- RQ1 Comparison of Commercial and Open-Source Models** *How do commercial LLMs compare with open-weight models on Burmese language tasks?*
- RQ2 Effect of Model Scale** *Does increasing model size lead to improved performance on Burmese NLP tasks?*
- RQ3 Effect of Southeast Asian (SEA) Fine-Tuning** *Does fine-tuning on Southeast Asian languages improve model performance on Burmese?*
- RQ4 Effect of Model Quantization** *How does model quantization affect performance compared to full-precision models?*

**RQ5 Temporal Progress in Burmese Language Capability** *Have LLMs demonstrated consistent improvements in Burmese performance across model generations?*

**Contributions.** We summarize the contribution of our work as follows:

- We propose *BURMESE-SAN*. A holistic Burmese benchmark for evaluating LLMs across seven NLP tasks. Our benchmark is formulated and edited by humans, resulting in high-quality samples.
- We present a reproducible dataset development process, covering sourcing, task design, and quality checks. <https://github.com/aisingapore/SEA-HELM>.
- We conduct extensive experiments on Instruction-Tuned, and Reasoning LLMs of different sizes, revealing systematic gaps in Burmese understanding, reasoning, and generation. The results show that SEA-tuned models can substantially improve performance.

## 2. Related Works

### 2.1. Burmese NLP

Myanmar is a linguistically diverse country with around 100 ethnic languages and dialects, belonging to four major language families: Sino-Tibetan, Austro-Asiatic, Tai-Kadai, and Indo-European (Ethnologue, 2024). Although Standard Burmese is the official language, several dialects are spoken, including Beik, Dawei, and Rakhine (Oo et al., 2023).

<sup>1</sup>The word *San* means *Standard* in Burmese.

Burmese (or Bamar), also known as Standard Burmese, is the native language of the Bamar majority and is also spoken by related groups such as the Mon. It belongs to the Sino-Tibetan (specifically Tibeto-Burman) language family. Burmese has been significantly influenced by Pali, the liturgical language of Theravada Buddhism, as well as by Mon and English. Over time, many foreign words entered Burmese as loanwords. After the end of British colonial rule, the government replaced many English terms by creating new Burmese equivalents (MULTICSD Project Team, 2025). Although Standard Burmese is the official language of Myanmar, widely used by the majority population and more resource-rich than other ethnic languages, it is still considered under-resourced for NLP. It lacks sufficient high-quality data, tools, and benchmarks for effective language processing. However, due to its official status, widespread use, and relative resource availability, we focus on Standard Burmese for benchmarking in this work.

The design of *BURMESE-SAN* focuses on linguistic authenticity and representativeness in vocabulary, sentence structure, and grammar. This enables a comprehensive evaluation of Natural Language Processing (NLP) capabilities in Burmese, encompassing Natural Language Understanding (NLU), Natural Language Reasoning (NLR), and Natural Language Generation (NLG). Regarding sentence structure, *BURMESE-SAN* is designed to reflect natural and native use of Standard Burmese. We carefully prepared the evaluation data to help evaluate how well the LLMs handle real-world Burmese language use. (See Section 3). The benchmark uses commonly spoken vocabulary, focusing on words and phrases frequently used in daily conversation. It also naturally includes loanwords - mainly from Pali and English - that are now a regular part of modern Burmese lexicon.

## 2.2. Dataset and Benchmark Overview

Previous efforts to evaluate Burmese language models have been mostly scattered and uncoordinated. Researchers have developed datasets for Burmese NLP tasks, including text classification, sequence tagging, and machine translation (Thu et al., 2023; San et al., 2024; Zaw et al., 2022; Hlaing et al., 2022; Aung et al., 2024; Kyaw et al., 2024). However, there is still no unified, comprehensive benchmark to systematically evaluate LLMs across different linguistic aspects of the Burmese language.

Recent Southeast Asian benchmarks largely exclude Burmese or rely on machine-translated data. SeaExam and SeaBench (Liu et al., 2025) introduce localized exam-style and open-ended queries

in Vietnamese, Thai, and Indonesian, but do not include Burmese. SailCompass (Guo et al., 2024b) evaluates Southeast Asian language understanding without any Burmese components. SeaEval (Wang et al., 2024) provides a multilingual and multicultural benchmark comprising 29 datasets; however, Burmese is excluded, and most of the data are machine-translated or derived from general-purpose multilingual prompts. BHASA covers Indonesian, Tamil, Thai, and Vietnamese (Leong et al., 2023), whereas IndoNLU focuses solely on Indonesian NLU tasks (Wilie et al., 2020), lacking evaluations of generative or reasoning capabilities.

Unlike these existing benchmarks, *BURMESE-SAN* provides a **comprehensive, multi-task evaluation suite for Burmese NLP**, covering NLU, NLR, and NLG tasks. All datasets were carefully curated and adapted to ensure native speaker verification, providing linguistic authenticity, high-quality labels, and cultural relevance. This makes *BURMESE-SAN* the **first holistic benchmark specifically designed for Burmese**.

## 3. Task and Dataset Curation

As shown in Table 2, *BURMESE-SAN* covers tasks across Natural Language Understanding (NLU), Natural Language Reasoning (NLR), and Natural Language Generation (NLG) with balanced class distributions, totaling 3,920 samples. We discuss how to formulate *BURMESE-SAN*, starting from task and dataset selection in Sections 3.1 and 3.2, to ensure that our benchmark is holistic. Finally, we then describe the dataset adaptation and annotation process in Section 3.3, including how we ensure the quality of our benchmark. Note that more information about the tasks and datasets, along with task descriptions, dataset information, and quality, is discussed further in Appendix A, B, and C.

### 3.1. Task Selection

Similar to well-exhibited benchmarks in the SEA languages (Montalan et al., 2025; Liu et al., 2025), we require tasks that are widely used to evaluate the robustness of LLMs that reflect real-world scenarios. The tasks should consist of classical NLP tasks, such as classification, and LLM-widely adopted tasks, including question answering (QA), summarization (AS), and machine translation (MT), as well as reasoning and understanding tasks, including natural language inference (NLI). Therefore, *BURMESE-SAN* brings comprehensive LLM evaluation to Burmese with seven tasks across NLU (sentiment analysis, question answering, and toxicity detection), NLG (abstractive summarization and machine translation), and NLR (causal reason-

Comp.	Task	Dataset	Target <sup>1</sup>	Language	Source <sup>2</sup>	Adaptation	Our contribution
NLU	QA	Belebele (Bandarkar et al., 2024)	S (4)	Burmese	English-adapted	Native re-translation	✓
	SA	GKLMIP-mya (Jiang et al., 2021)	L (3)	Burmese	Natively-sourced	Normalization <sup>3</sup>	✓
	TD	myHateSpeech (Kyaw et al., 2024)	L (9)	Burmese	Natively-sourced	No Adaptation	
NLR	CR	Balanced COPA (Kavumba et al., 2019)	L (2)	English	English-sourced	Native translation	✓
	NLI	myXNLI (Htet and Dras, 2025)	L (3)	Burmese	Natively-translated	No Adaptation	
NLG	AS	XL-Sum (Hasan et al., 2021)	Sum	Burmese	English-adapted	Native re-translation	✓
	MT	FLORES+ (Team, 2024a)	Trans	Burmese	English-adapted	Native re-translation	✓

<sup>1</sup> Number of options shown in parentheses. L: Label; S: Span; Sum: Summary; Trans: Translate

<sup>2</sup> **English-adapted**: previously translated from English, then further adapted or corrected by native speakers.

**English-sourced**: originally written in English.

**Natively-sourced**: originally created in Burmese by native speakers. **Natively-translated**: translated by native speakers

<sup>3</sup> **Normalization**: spelling was standardized to follow consistent orthography and modern usage norms, including Unicode normalization and correction of common typos.

Table 1: Dataset information for each task in *BURMESE-SAN*. A checkmark in the *Our contribution* column indicates direct contributions to the adaptation process.

Competency	Task	Label	# samples
NLU	QA	Span	120
		SA	Positive
		Negative	200
		Neutral	200
	TD	Clean	200
		Hate	200
NLR	CR	Cause	200
		Effect	200
	NLI	Contradiction	200
		Entailment	200
		Neutral	200
NLG	AS	Summary	100
	MT	MYA → ENG	850
		ENG → MYA	850
Total			3,920

Table 2: Class distribution per task in *BURMESE-SAN* Benchmark Dataset

ing and natural language inference), as shown in Table 1. This ensures that our benchmark covers all aspects of LLM evaluation.

### 3.2. Dataset Selection

To build *BURMESE-SAN*, we collected open-source datasets with clear sources, focusing on those that reflect authentic Burmese language use in domains such as social media, news, and forums. For tasks like TD and NLI, high-quality datasets are already available that do not require adaptation. These datasets are either formulated in Burmese (monolingual datasets) or carefully edited by previous works to fix translation issues (Kyaw et al., 2024; Htet and Dras, 2025).

However, for other tasks, existing datasets required translation or editing to meet our goals of creating high-quality and culturally relevant datasets. Therefore, we adapted English and partially translated data sets, including FLORES+ (Team, 2024a), Belebele (Bandarkar et al., 2024), GKLMIP-mya (Jiang et al., 2021), XL-Sum (Hasan et al., 2021), and Balanced COPA (Kavumba et al., 2019), to create Burmese-SAN. While these datasets are

well-known and widely used in LLM evaluation benchmarks (Lovenia et al., 2024; Susanto et al., 2025), they were not readily available in Burmese or were inadequately translated. We applied normalization and native re-translation to ensure suitability for inclusion in Burmese-SAN. Table 1 also summarizes the usage, adaptations, and our contributions, highlighting the *comprehensive, native-speaker-verified nature of Burmese-SAN*, which distinguishes it from prior benchmarks (Vayani et al., 2025; Lovenia et al., 2024).

### 3.3. Dataset Adaptation and Annotation

As we discussed in the previous discussion, although some Burmese datasets are available, their data format and quality need refinement to make them *natural in Burmese, culturally suitable, and reliable for evaluation*. As shown in Figure 1 (Right), we designed a four-step process (random sampling and filtering, translation, text normalization, and label verification, as well as native revision) to combine Burmese and non-Burmese datasets into a single benchmark. Note that, to ensure dataset quality, we employed native Burmese speakers (ages 18–25), primarily university students enrolled in international programs as annotators.

**Random Sampling and Filtering** For each task, we first sampled data and restricted text length to 20–3,000 characters, ensuring balanced class distributions to reduce bias. All NLU and NLR tasks were limited to 32 tokens per output, whereas generation tasks utilized longer outputs, with 256 tokens for MT and 512 tokens for AS, to accommodate richer content generation. Low-quality items, such as duplicates, unclear text, or unnatural writing, were removed by native speakers to maintain the high quality of Burmese examples. For the natively sourced datasets (SA, TD) and the natively translated NLI (Htet and Dras, 2025), we applied task-specific filtering. For the SA and TD tasks, we removed sentences with contextually ambiguous expressions that could obscure labels. For the

NLI task, we ensured premise–hypothesis pairs accurately reflected the intended relation.

**Translation and Binary Criteria Rating** For tasks originally in English (QA, CR, AS, MT), the data was translated into Burmese using annotators, ensuring high-quality data. For the English-sourced dataset (CR), we translated it from English to Burmese by leveraging our native speaker annotators. In contrast, the English-adapted datasets (QA, AS, and MT) were re-translated due to the low quality of existing translations. After translation, our native-speaker annotators applied a binary rating (agree/disagree) to evaluate dataset quality based on multiple criteria. We used Completeness, Fluency, and Sensibility to assess information retention, grammaticality, and contextual logic for all tasks; Faithfulness for translation to evaluate meaning preservation; and Relevance and Coherence for summarization to measure content focus and structural flow. The initial overall mean joint agreement across all tasks and criteria is 99.11%, indicating a consistently high level of annotator alignment throughout the evaluation. For each sample, if it did not receive full agreement from our annotators, it was then revised by them to improve quality. This process efficiently filtered out poor-quality translations without requiring full re-annotation, ensuring that only accurate and fluent Burmese text was included in the benchmark.

**Text Normalization and Label Verification** In the next step, normalization was applied to improve consistency for datasets that were natively sourced (SA and TD). This process primarily served to resolve ambiguities and measure inter-annotator agreement, preserving dataset integrity while improving reliability for downstream evaluation. For Toxicity Detection (TD), despite spelling errors, distorted spellings were preserved, as they reflect real social media usage and intentional attempts to bypass detection. For Sentiment Analysis (SA), typos, spelling mistakes, and character-order issues were corrected so that errors would not affect classification. Moreover, code-switching (English-Burmese) samples were preserved in both tasks, as it is a natural part of Myanmar social media. For the gold label of each sample, due to the subjective nature of these tasks, native speakers verified and, when necessary, adjusted the labels. We retained the original labels as far as possible, as our goal was to ensure label consistency and quality rather than redefine the task. We found that the agreement between the original and verified labels was very high. For Sentiment Analysis (SA), Cohen’s Kappa = 0.948, Krippendorff’s Alpha = 0.948, and total agreement = 96.5%. For Toxicity Detection (TD), Cohen’s Kappa = 0.985, Krippendorff’s Alpha = 0.985, and total agreement = 99.25%.

**Native Revision** Finally, native speakers con-

ducted a final review of all text. This step focused on refining linguistic naturalness, cultural appropriateness, and readability, including minor adjustments to grammar, punctuation, and style. The goal was not to alter the meaning, but to ensure that the dataset reflects authentic Burmese language use suitable for downstream tasks.

## 4. Experimental Setup

### 4.1. Prompt Templates

For *BURMESE-SAN*, we designed task-specific prompt templates entirely in Burmese. These templates were carefully aligned with the principles of prompt design established in SEA-HELM (Suantano et al., 2025) to ensure consistency between tasks. In particular, we translate the task prompts from SEA-HELM into the native Burmese language. This is because the design prompt of SEA-HELM maintains a clear separation between instruction and content, as the formal prompts are unambiguous and standardized, thereby reducing confounding factors that might influence model performance due to differences in prompt style rather than genuine understanding of the Burmese text.

### 4.2. Evaluation Setup

**Tasks.** Each task in *BURMESE-SAN*, as described in Section 3, is designed to thoroughly test how well LLMs understand and use the Burmese language. A task includes a set of test and example cases, each with an input, reference answers, metadata, and the correct label. **Metrics** We adopt multiple evaluation metrics to assess model performance across tasks. For each metric, the model output is treated as the prediction, while the corresponding instance label serves as the reference. For NLU and NLR tasks, we report the accuracy score. For machine translation, we use MetricX-24 (Juraska et al., 2024). For abstractive summarization, we report ROUGE-L F1 using the multilingual ROUGE implementation from XL-Sum (Hasan et al., 2021). All metric scores are normalized to a common [0,100] scale following the SEA-HELM normalization process, which accounts for differences in metric ranges, task difficulty, and random baselines (e.g., MetricX originally ranges from [0,25] but is rescaled to [0,100]). Each model is evaluated across eight independent runs without greedy decoding (i.e., temperature = 0) (Miller, 2024), and we report the mean normalized score to ensure stable and reliable performance estimates.

### 4.3. Models

**Model Selection.** We evaluated a set of LLMs across diverse families, parameter scales, and ac-

Model	MY	CR	NLI	QA	SA	TD	AS	MT
<i>Small Models (&lt; 14B)</i>								
MERaLION 2 (10B)	10.66 ± 0.16	0.00 ± 0.00	12.61 ± 0.31	0.00 ± 0.00	0.07 ± 0.13	14.79 ± 0.09	19.49 ± 0.27	35.33 ± 0.15
Olmo 2 1124 (13B)	1.84 ± 0.09	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.79 ± 0.03	0.00 ± 0.00	1.04 ± 0.03
Olmo 2 1124 (7B)	2.18 ± 0.12	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	1.80 ± 0.04	0.00 ± 0.00	6.44 ± 0.10
Olmo 3 (7B)	3.41 ± 0.11	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	3.66 ± 0.05	14.00 ± 0.14	1.11 ± 0.02
Tulu 3 (8B)	10.95 ± 0.14	0.00 ± 0.00	1.35 ± 0.13	5.89 ± 0.64	0.00 ± 0.00	13.70 ± 0.06	<u>21.48 ± 0.17</u>	37.48 ± 0.10
SEA-LION v3-Gemma-2 (9B)	15.40 ± 0.22	0.00 ± 0.00	25.10 ± 0.37	4.52 ± 1.35	0.00 ± 0.00	20.64 ± 0.11	14.84 ± 0.29	53.10 ± 0.15
SEA-LION v3-Llama (8B)	17.88 ± 0.26	4.45 ± 1.29	17.01 ± 0.70	37.41 ± 0.81	22.19 ± 0.97	12.79 ± 0.26	5.61 ± 0.27	13.92 ± 0.12
SEA-LION v4-Apertus (8B)	16.68 ± 0.22	0.00 ± 0.00	0.00 ± 0.00	23.44 ± 0.97	19.05 ± 0.43	15.71 ± 0.08	5.47 ± 0.13	55.46 ± 0.16
SEA-LION v4-Gemma-3-VL (4B)	26.24 ± 0.13	0.00 ± 0.00	<u>34.39 ± 0.30</u>	47.78 ± 0.00	<u>43.11 ± 0.15</u>	27.90 ± 0.07	17.93 ± 0.10	<u>62.56 ± 0.12</u>
SEA-LION v4-Qwen-3-VL (4B)	23.31 ± 0.17	28.80 ± 0.41	21.26 ± 0.21	42.81 ± 0.25	31.54 ± 0.14	20.37 ± 0.10	0.63 ± 0.03	24.90 ± 0.13
SEA-LION v4-Qwen-3-VL (8B)	<u>30.67 ± 0.19</u>	<u>38.77 ± 0.41</u>	23.28 ± 0.25	<u>58.22 ± 0.60</u>	36.21 ± 0.27	31.76 ± 0.11	15.49 ± 0.21	51.70 ± 0.20
Qwen 2.5 (7B)	8.04 ± 0.15	0.00 ± 0.00	5.97 ± 0.82	0.00 ± 0.00	0.00 ± 0.00	8.77 ± 0.15	19.35 ± 0.18	11.22 ± 0.11
Qwen 3 VL (4B)	18.61 ± 0.13	0.00 ± 0.00	29.18 ± 0.48	40.59 ± 0.41	35.60 ± 0.13	21.34 ± 0.09	21.48 ± 0.09	40.17 ± 0.12
Qwen 3 VL (8B)	26.26 ± 0.20	30.70 ± 0.65	3.16 ± 0.14	40.04 ± 0.73	33.83 ± 0.33	25.73 ± 0.13	20.11 ± 0.15	47.90 ± 0.13
Qwen 3 VL (4B)	20.42 ± 0.16	22.77 ± 0.57	21.53 ± 0.32	44.44 ± 0.38	36.70 ± 0.13	20.01 ± 0.13	0.55 ± 0.00	34.34 ± 0.11
Qwen 3 VL (8B)	30.25 ± 0.22	36.85 ± 0.59	32.05 ± 0.46	53.59 ± 0.70	36.41 ± 0.26	<u>32.60 ± 0.17</u>	16.51 ± 0.23	49.72 ± 0.09
Babel (9B)	8.01 ± 0.18	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	22.39 ± 0.63	8.64 ± 0.08	18.33 ± 0.33	17.51 ± 0.22
SeaLLMs V3 (7B)	7.13 ± 0.16	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	8.47 ± 0.08	17.68 ± 0.31	18.87 ± 0.15
Aya Expanse (8B)	3.03 ± 0.13	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	1.48 ± 0.04	0.00 ± 0.00	2.82 ± 0.06
Command R7B 12-2024 (7B)	3.19 ± 0.13	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	4.36 ± 0.07	10.06 ± 0.28	9.73 ± 0.14
Gemma 2 (9B)	9.63 ± 0.17	0.00 ± 0.00	9.16 ± 0.63	0.00 ± 0.00	0.00 ± 0.00	13.02 ± 0.13	14.08 ± 0.28	35.64 ± 0.13
Gemma 3 VL (12B)	<b>42.46 ± 0.15</b>	<b>44.93 ± 0.77</b>	26.22 ± 0.16	<b>67.11 ± 0.41</b>	<b>48.96 ± 0.22</b>	<b>40.97 ± 0.15</b>	18.77 ± 0.26	<b>70.96 ± 0.09</b>
Gemma 3 VL (4B)	20.56 ± 0.19	0.00 ± 0.00	22.12 ± 0.31	35.74 ± 0.38	39.18 ± 0.16	22.25 ± 0.09	19.35 ± 0.12	50.90 ± 0.14
Llama 3 (8B)	3.20 ± 0.06	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	4.89 ± 0.03	0.00 ± 0.00	22.95 ± 0.15
Llama 3.1 (8B)	8.45 ± 0.19	0.00 ± 0.00	0.00 ± 0.00	7.48 ± 1.02	0.00 ± 0.00	10.72 ± 0.09	<b>26.73 ± 0.27</b>	20.69 ± 0.14
Ministral 2410 (8B)	3.90 ± 0.13	0.00 ± 0.00	0.29 ± 0.23	0.00 ± 0.00	0.05 ± 0.10	6.61 ± 0.07	4.22 ± 0.19	27.35 ± 0.17
Sailor2 (8B)	11.65 ± 0.13	0.00 ± 0.00	<b>43.74 ± 0.79</b>	0.26 ± 0.27	0.00 ± 0.00	19.30 ± 0.14	0.00 ± 0.00	48.76 ± 0.14
Apertus (8B)	9.30 ± 0.21	0.00 ± 0.00	0.25 ± 0.22	3.78 ± 1.27	2.92 ± 0.74	12.13 ± 0.10	13.22 ± 0.27	40.68 ± 0.16
<i>Medium Models (14B–32B)</i>								
Olmo 2 0325 (32B)	4.38 ± 0.15	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	3.91 ± 0.05	0.00 ± 0.00	14.73 ± 0.12
SEA-LION v4-Gemma-3-VL (27B)	47.18 ± 0.15	57.15 ± 0.47	52.92 ± 0.18	69.11 ± 0.41	<b>47.56 ± 0.26</b>	48.95 ± 0.11	11.43 ± 0.23	<u>77.83 ± 0.10</u>
SEA-LION v4-Qwen-3-VL (32B)	<b>49.56 ± 0.14</b>	<b>66.38 ± 0.20</b>	<b>57.13 ± 0.20</b>	<b>81.52 ± 0.22</b>	40.49 ± 0.13	<b>51.70 ± 0.08</b>	23.56 ± 0.08	64.00 ± 0.19
Qwen 2.5 (14B)	13.59 ± 0.19	0.00 ± 0.00	4.36 ± 0.68	33.96 ± 0.68	0.00 ± 0.00	14.25 ± 0.13	21.50 ± 0.11	32.46 ± 0.10
Qwen 2.5 (32B)	26.99 ± 0.17	13.68 ± 0.64	5.60 ± 0.47	43.04 ± 0.71	32.36 ± 0.33	22.78 ± 0.14	<u>24.05 ± 0.12</u>	39.36 ± 0.15
Qwen 3 VL (14B)	32.46 ± 0.13	31.75 ± 0.52	17.80 ± 0.30	61.26 ± 0.63	36.96 ± 0.23	31.52 ± 0.12	22.76 ± 0.12	51.87 ± 0.16
Qwen 3 VL (32B)	44.60 ± 0.19	<u>60.68 ± 0.45</u>	52.07 ± 0.34	<u>74.48 ± 0.73</u>	41.37 ± 0.31	45.33 ± 0.12	<b>25.31 ± 0.18</b>	44.71 ± 0.18
Qwen 3 VL (32B)	40.89 ± 0.19	49.48 ± 0.48	43.49 ± 0.28	74.04 ± 0.27	<u>47.17 ± 0.17</u>	40.26 ± 0.12	10.77 ± 0.23	56.03 ± 0.11
Aya Expanse (32B)	6.44 ± 0.15	0.25 ± 0.24	0.25 ± 0.22	0.00 ± 0.00	0.00 ± 0.00	5.67 ± 0.08	0.00 ± 0.00	20.66 ± 0.18
Command R 08-2024 (32B)	4.87 ± 0.15	0.00 ± 0.00	1.37 ± 0.69	1.07 ± 0.61	0.00 ± 0.00	6.30 ± 0.13	16.97 ± 0.15	9.75 ± 0.11
Gemma 2 (27B)	23.97 ± 0.24	25.78 ± 1.05	32.10 ± 0.27	31.30 ± 1.50	21.46 ± 0.87	29.59 ± 0.21	21.41 ± 0.19	50.32 ± 0.12
Gemma 3 VL (27B)	<u>48.14 ± 0.17</u>	57.08 ± 0.51	<u>53.82 ± 0.16</u>	69.85 ± 0.31	47.13 ± 0.24	<u>49.96 ± 0.12</u>	13.16 ± 0.25	<b>79.44 ± 0.07</b>
phi-4 (14B)	6.45 ± 0.18	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	17.77 ± 0.69	6.65 ± 0.09	10.79 ± 0.39	16.19 ± 0.12
Mistral Small 3.1 2503 (24B)	2.22 ± 0.11	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	2.38 ± 0.05	3.48 ± 0.09	6.34 ± 0.14
Sailor2 (20B)	8.55 ± 0.11	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	11.66 ± 0.04	0.00 ± 0.00	52.86 ± 0.15
<i>Large Models (&gt; 32B)</i>								
Tulu 3 (70B)	35.11 ± 0.17	56.78 ± 0.51	43.75 ± 0.43	66.85 ± 0.68	34.50 ± 0.40	43.89 ± 0.14	25.97 ± 0.24	66.60 ± 0.09
SEA-LION v3-Llama (70B)	38.21 ± 0.28	54.30 ± 0.91	44.63 ± 0.47	73.30 ± 0.70	0.00 ± 0.00	43.98 ± 0.20	25.26 ± 0.16	63.28 ± 0.16
Qwen 3 A3B (30B MoE)	25.62 ± 0.12	0.00 ± 0.00	44.07 ± 0.47	69.04 ± 0.44	31.12 ± 0.29	21.90 ± 0.09	1.27 ± 0.12	34.84 ± 0.12
Qwen 2.5 (72B)	27.54 ± 0.20	5.58 ± 0.22	12.41 ± 0.25	48.63 ± 0.45	39.85 ± 0.22	23.51 ± 0.10	21.68 ± 0.27	46.32 ± 0.14
Qwen 3 A22B (235B MoE)	<u>54.29 ± 0.16</u>	<b>71.32 ± 0.23</b>	<b>57.42 ± 0.19</b>	76.63 ± 0.07	<b>50.32 ± 0.18</b>	<b>56.51 ± 0.08</b>	23.36 ± 0.14	78.40 ± 0.06
Qwen 3 Next (80B MoE)	44.88 ± 0.16	59.27 ± 0.60	46.77 ± 0.22	74.89 ± 0.43	41.88 ± 0.20	46.01 ± 0.12	21.66 ± 0.08	58.58 ± 0.15
Babel (83B)	9.87 ± 0.23	0.00 ± 0.00	0.40 ± 0.32	29.89 ± 2.08	8.11 ± 1.02	7.72 ± 0.10	16.56 ± 0.24	9.61 ± 0.15
ERNIE 4.5 (21B MoE)	17.70 ± 0.15	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	19.67 ± 0.06	10.37 ± 0.19	72.24 ± 0.18
ERNIE 4.5 (300B MoE)	<b>54.68 ± 0.16</b>	68.82 ± 0.22	40.39 ± 0.15	<u>78.52 ± 0.32</u>	<u>49.81 ± 0.16</u>	<u>54.27 ± 0.08</u>	20.84 ± 0.22	<b>86.25 ± 0.05</b>
Command A 03-2025 (111B)	16.52 ± 0.23	0.00 ± 0.00	18.57 ± 0.58	30.78 ± 1.22	16.68 ± 0.75	15.71 ± 0.14	5.57 ± 0.35	37.10 ± 0.19
Command R+ 08-2024 (104B)	6.61 ± 0.18	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	1.83 ± 0.68	9.19 ± 0.07	16.00 ± 0.17	25.95 ± 0.17
DeepSeek V3 (671B MoE)	40.87 ± 0.25	59.22 ± 0.80	16.31 ± 0.79	15.78 ± 1.72	39.07 ± 0.33	41.73 ± 0.17	20.22 ± 0.13	72.92 ± 0.13
DeepSeek V3.1 (671B MoE)	51.30 ± 0.23	58.92 ± 0.70	<u>51.64 ± 0.37</u>	70.41 ± 0.60	44.17 ± 0.21	53.41 ± 0.16	21.47 ± 0.11	<u>85.86 ± 0.05</u>
Llama 3 (70B)	13.09 ± 0.17	32.92 ± 0.66	35.18 ± 0.43	15.33 ± 1.12	0.00 ± 0.00	24.03 ± 0.14	0.00 ± 0.00	49.89 ± 0.11
Llama 3.1 (70B)	19.87 ± 0.24	0.00 ± 0.00	35.08 ± 0.71	37.63 ± 1.30	0.00 ± 0.00	26.76 ± 0.15	27.30 ± 0.26	58.42 ± 0.20
Llama 3.3 (70B)	23.07 ± 0.15	0.00 ± 0.00	42.51 ± 0.31	3.70 ± 0.72	0.00 ± 0.00	29.50 ± 0.09	<u>28.32 ± 0.21</u>	60.06 ± 0.15
Llama 4 Maverick (400B MoE)	51.49 ± 0.20	67.87 ± 0.21	31.87 ± 0.39	<b>80.00 ± 0.00</b>	44.92 ± 0.11	52.54 ± 0.10	<b>30.66 ± 0.15</b>	81.86 ± 0.05
Llama 4 Scout (109B MoE)	45.54 ± 0.17	<u>70.40 ± 0.32</u>	31.74 ± 0.21	75.93 ± 0.19	13.19 ± 0.43	49.98 ± 0.09	25.53 ± 0.25	81.15 ± 0.04
Mistral Large 2411 (123B)	26.17 ± 0.29	7.00 ± 1.42	10.76 ± 0.40	41.41 ± 1.39	36.23 ± 0.41	24.14 ± 0.27	19.68 ± 0.28	55.08 ± 0.20
Kimi K2 Instruct 0905 (1040B MoE)	43.94 ± 0.19	53.02 ± 1.07	13.98 ± 0.36	58.81 ± 0.94	40.21 ± 0.47	41.64 ± 0.20	23.02 ± 0.20	71.96 ± 0.19
Apertus (70B)	13.09 ± 0.25	0.00 ± 0.00	0.00 ± 0.00	21.78 ± 1.43	0.00 ± 0.00	16.60 ± 0.09	15.17 ± 0.17	58.27 ± 0.17

Table 3: Performance of instruct models on BURMESE-SAN tasks. Best model for each task per size group is bold, second best is underlined.

cess types. The selection includes (i) instruction-tuned models such as Qwen 3 (Yang and et al., 2025; Bai and et al., 2025), Llama-3/3.1/3.3/4 series (Grattafiori and et al., 2024), Gemma 2/3 (Team, 2025c), SEA-LION (Ng et al., 2025), DeepSeek-V3, V3.1 (671B MoE) (Team, 2025a), OLMo (Groeneveld et al., 2024), Mistral (Jiang et al., 2023), Sailor2-Chat (Dou et al., 2025), Babel (Zhao et al., 2025), and Command-R series; (ii) reasoning-focused models such as DeepSeek-

V3.1 Thinking and DeepSeek-R1 (DeepSeek-AI, 2025), Qwen3-Thinking (Yang and et al., 2025), QwQ (Team, 2025f), GPT-OSS (Team, 2025d), OLMo-3-Think (Olmo, 2025), and SEA-LION v3.5 R (Ng et al., 2025); and (iii) commercial models including Google’s Gemini-2/2.5 (Flash and Pro) (Team, 2025b), OpenAI’s GPT-4o (Team, 2024b), GPT-4.1 (OpenAI, 2024), and GPT-5 (Team., 2025), and Anthropic’s Claude family (Sonnet 4 and Opus 4.1) (Anthropic, 2025). This diverse collection ensures

Model	MY	CR	NLI	QA	SA	TD	AS	MT
<i>Small Models (&lt; 14B)</i>								
Olmo 3 Think (7B)	3.85 ± 0.16	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	5.88 ± 0.06	9.99 ± 0.23	17.61 ± 0.12
SEA-LION v3.5 R (Llama) (8B)	22.08 ± 0.20	0.00 ± 0.00	30.19 ± 0.64	54.00 ± 1.22	0.00 ± 0.00	22.28 ± 0.15	11.02 ± 0.32	48.28 ± 0.32
Qwen 3 (Thinking) (4B)	34.78 ± 0.29	40.78 ± 1.30	31.98 ± 0.70	60.37 ± 0.96	31.82 ± 0.53	35.77 ± 0.28	16.29 ± 0.40	56.00 ± 0.11
Qwen 3 (Thinking) (8B)	<b>37.12 ± 0.22</b>	<b>55.80 ± 0.95</b>	<b>33.45 ± 0.63</b>	<b>65.19 ± 1.23</b>	<b>40.18 ± 0.37</b>	<b>40.22 ± 0.22</b>	<b>18.28 ± 0.18</b>	<b>59.53 ± 0.12</b>
<i>Medium Models (14B–32B)</i>								
Olmo 3 Think (32B)	5.51 ± 0.16	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	5.85 ± 0.07	8.80 ± 0.19	15.28 ± 0.14
QwQ (32B)	24.66 ± 0.27	0.00 ± 0.00	19.04 ± 0.69	55.96 ± 1.33	30.38 ± 0.49	18.58 ± 0.15	4.17 ± 0.27	38.97 ± 0.19
Qwen 3 (Thinking) (14B)	<u>43.81 ± 0.24</u>	<u>61.45 ± 0.97</u>	35.47 ± 0.45	<u>73.70 ± 1.26</u>	<u>38.71 ± 0.35</u>	<u>45.07 ± 0.20</u>	<u>20.86 ± 0.12</u>	<b>65.04 ± 0.21</b>
Qwen 3 (Thinking) (32B)	<b>52.35 ± 0.26</b>	<b>71.32 ± 0.70</b>	<b>48.61 ± 0.54</b>	<b>77.33 ± 0.87</b>	<b>42.32 ± 0.35</b>	<b>48.49 ± 0.19</b>	<b>22.90 ± 0.16</b>	43.43 ± 0.32
Reka Flash 3.1 (21B)	24.51 ± 0.28	3.80 ± 1.34	<u>39.16 ± 0.65</u>	61.00 ± 1.34	5.20 ± 0.75	26.51 ± 0.28	10.91 ± 0.21	<u>56.18 ± 0.19</u>
<i>Large Models (&gt; 32B)</i>								
SEA-LION v3.5 R (Llama) (70B)	44.02 ± 0.26	19.65 ± 1.58	41.37 ± 0.56	75.07 ± 0.67	8.83 ± 0.80	41.60 ± 0.29	<b>21.55 ± 0.10</b>	78.97 ± 0.13
Qwen 3 (Thinking) (235B MoE)	<b>66.39 ± 0.22</b>	<b>77.97 ± 0.44</b>	<b>51.53 ± 0.47</b>	<b>82.44 ± 0.58</b>	<b>44.68 ± 0.26</b>	<b>61.37 ± 0.13</b>	20.48 ± 0.09	85.48 ± 0.03
Qwen 3 (Thinking) (30B MoE)	53.54 ± 0.19	66.55 ± 0.77	48.84 ± 0.59	76.26 ± 0.78	38.92 ± 0.27	53.77 ± 0.16	<u>21.41 ± 0.15</u>	78.70 ± 0.08
Qwen 3 Next (Thinking) (80B MoE)	57.97 ± 0.18	72.50 ± 0.52	<u>49.68 ± 0.54</u>	<u>80.04 ± 0.66</u>	40.75 ± 0.30	56.35 ± 0.15	20.08 ± 0.12	79.88 ± 0.08
DeepSeek V3.1 Thinking (671B MoE)	<u>59.46 ± 0.23</u>	71.65 ± 0.67	48.27 ± 0.58	78.26 ± 0.86	40.91 ± 0.32	<u>57.75 ± 0.15</u>	21.21 ± 0.13	<b>86.46 ± 0.04</b>
Deepseek R1 0528 (671B MoE)	57.31 ± 0.30	<u>72.57 ± 0.54</u>	47.25 ± 0.70	66.07 ± 1.29	<u>41.44 ± 0.31</u>	56.41 ± 0.17	18.28 ± 0.19	<u>85.71 ± 0.05</u>
GPT OSS (120B MoE mxfp4)	48.71 ± 0.21	58.95 ± 0.78	24.25 ± 0.53	70.19 ± 0.89	39.11 ± 0.43	46.00 ± 0.16	17.34 ± 0.11	78.03 ± 0.08
GPT OSS (20B MoE mxfp4)	35.25 ± 0.24	37.58 ± 1.00	13.97 ± 0.93	60.07 ± 1.37	36.12 ± 0.44	34.42 ± 0.26	17.17 ± 0.16	67.28 ± 0.09

Table 4: Performance of Reasoning models on BURMESE-SAN tasks. Best model for each task per size group is bold, second best is underlined.

Model	MY	CR	NLI	QA	SA	TD	AS	MT
Opus 4.1 (2025-08-05)	45.10 ± 0.15	71.57 ± 0.35	0.00 ± 0.00	0.00 ± 0.00	48.35 ± 0.19	45.66 ± 0.08	<u>24.69 ± 0.09</u>	87.52 ± 0.03
Sonnet 4	40.45 ± 0.30	24.33 ± 0.13	<b>86.01 ± 0.03</b>	0.00 ± 0.00	47.82 ± 0.25	43.22 ± 0.15	<b>68.10 ± 0.67</b>	0.00 ± 0.00
Gemini 2 Flash	59.23 ± 0.16	76.58 ± 0.43	41.30 ± 0.43	79.81 ± 0.33	41.48 ± 0.24	57.78 ± 0.12	22.09 ± 0.09	88.23 ± 0.04
Gemini 2.5 Flash	<u>69.34 ± 0.16</u>	<u>83.58 ± 0.40</u>	61.07 ± 0.46	<u>84.15 ± 0.53</u>	39.60 ± 0.21	<u>66.16 ± 0.10</u>	24.11 ± 0.14	<u>89.49 ± 0.02</u>
Gemini 2.5 Pro	<b>72.35 ± 0.18</b>	<b>85.75 ± 0.38</b>	<u>67.55 ± 0.43</u>	<b>85.48 ± 0.53</b>	45.56 ± 0.30	<b>68.74 ± 0.11</b>	24.21 ± 0.12	<b>90.22 ± 0.02</b>
GPT 4.1 (2025-04-14)	55.80 ± 0.18	68.80 ± 0.37	42.80 ± 0.38	77.15 ± 0.31	<u>50.24 ± 0.27</u>	54.05 ± 0.11	21.37 ± 0.24	79.73 ± 0.08
GPT 4o (2024-11-20)	51.61 ± 0.20	70.77 ± 0.53	39.48 ± 0.43	75.67 ± 0.66	<b>50.94 ± 0.35</b>	52.34 ± 0.13	21.97 ± 0.13	78.58 ± 0.11
GPT 5 (2025-08-07)	66.46 ± 0.19	79.87 ± 0.66	43.63 ± 0.44	81.37 ± 0.47	45.70 ± 0.30	60.16 ± 0.13	17.09 ± 0.07	87.46 ± 0.04

Table 5: Performance of Commercial models on BURMESE-SAN tasks. Best model for each task per size group is bold, second best is underlined.

representative coverage of both open-source and proprietary ecosystems and enables analysis of scaling trends from small to large models.

**Inference Details.** During evaluation, input prompts are constructed by combining evaluation instances with their corresponding prompt templates. The default evaluation setting in BURMESE-SAN for instruction-tuned models is zero-shot prompting, where prompts do not include any in-context input-label examples. For decoding parameters, model-specific default configurations are used when available. For any unspecified parameters, we apply vLLM default settings. Given the input prompts and decoding parameters, the model generates outputs for evaluation.

## 5. Evaluation Results

We present our findings organized around five research questions that examine key aspects of Burmese language model capabilities as described in Section 1. Tables 3, 4, and 5 report performance across all NLP tasks, where the MY column denotes overall performance. Figure 2 compares original models with their SEA-fine-tuned variants (left) and SEA-LION models with their quantized versions (right), highlighting the effects of regional fine-tuning and quantization.

**Finding #1: Commercial models consistently outperform open-weight models (RQ1).** Com-

mercial models achieve substantially higher performance on Burmese tasks, led by Gemini 2.5 Pro (72.35%), Gemini 2.5 Flash (69.34%), and GPT-5 (66.46%). In contrast, the strongest open-weight models - ERNIE 4.5 (54.68%), Qwen 3 235B MoE (54.29%), and Llama 4 Maverick (51.49%) - lag behind, with a gap of approximately 17.67% between the top commercial and open-weight models. This disparity is especially pronounced in tasks requiring cultural and language-specific reasoning.

**Finding #2: Larger models tend to perform better, but scale alone is insufficient (RQ2).** Performance generally improves with model scale, though gains are non-linear and diminish at larger sizes. While larger variants within families such as Qwen 3 and Gemma 3 outperform smaller ones, notable exceptions exist: DeepSeek V3.1 substantially exceeds V3 (51.30% vs. 40.87%), and the much larger Kimi K2 Instruct (1040B MoE, 43.94%) underperforms smaller models. These results highlight the critical role of architecture, training data, and tuning strategies beyond parameter count.

**Finding #3: Southeast Asian fine-tuning benefits certain model families (RQ3).** SEA fine-tuning yields selective improvements that depend strongly on the base model. While Qwen-based SEA-LION v4 (32B) shows moderate gains (+4.96%) and the Gemma-based variant slightly degrades (-0.96%), Llama-based SEA-LION models benefit substantially from SEA fine-tuning, particu-

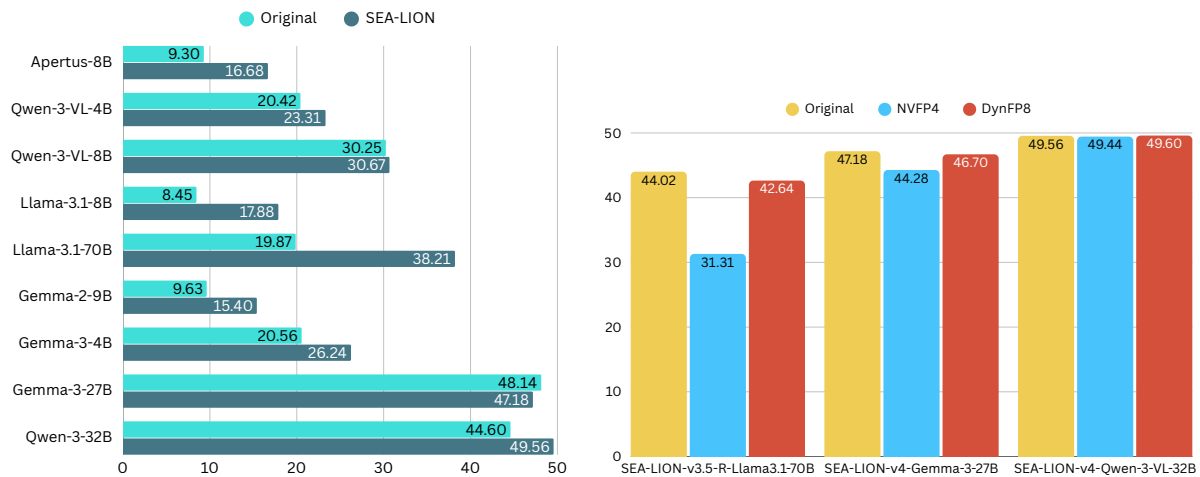


Figure 2: Left: Comparison of original models against SEA-fine-tuned variants, and Right: SEA-LION models with their quantized versions - NVIDIA FP4 (NVFP4) and Dynamic FP8 (DynFP8).

larly at larger scales. For example, SEA-LION v3 (Llama 3.1) 70B improves markedly over its base counterparts, indicating that SEA-specific data is especially effective for Llama architectures. Task-wise, gains are most pronounced in machine translation (+19.29%) and question answering (+7.04%), highlighting the value of regional fine-tuning for cross-lingual and generation-centric tasks.

**Finding #4: Careful quantization preserves performance for most tasks (RQ4).** Modern quantization methods can largely retain model performance when applied conservatively. For SEA-LION v4 (Qwen) 32B, both 8-bit DynFP8 and 4-bit NVFP4 quantization yield results comparable to full precision. DynFP8 similarly preserves performance for Gemma- and Llama-based models, whereas aggressive NVFP4 quantization causes notable degradation, particularly for reasoning-intensive models. These results indicate that quantization effectiveness depends on both the chosen method and the target task.

**Finding #5: Burmese language capability has improved rapidly across model generations (RQ5).** Burmese performance has increased substantially across successive model generations, with clear recent acceleration. Major open-weight families such as Llama, Qwen, and Gemma exhibit large generational gains (e.g., Llama 3.3 70B at 23.07% to Llama 4 Maverick at 51.49%), while commercial models show steady progress from GPT-4o (51.61%) to GPT-5 (66.46%) and from Gemini 2 Flash (59.23%) to Gemini 2.5 Pro (72.35%). Similar trends are observed within the SEA-LION series, where newer releases consistently outperform earlier versions; for instance, SEA-LION v4 (Gemma 3) 4B achieves 26.24%, substantially exceeding SEA-LION v3 (Gemma 2) 9B at 15.40%.

Commercial models continue to achieve the strongest performance, but the gap with open-weight models is steadily narrowing as architec-

tures and training strategies improve. While model scale remains relevant, performance depends on architectural design, data quality, and instruction tuning rather than parameter count alone.

Regional fine-tuning yields model-dependent benefits - particularly for Llama-based models. Finally, temporal analysis highlights rapid and consistent improvements in Burmese language capability across model generations, offering practical guidance for model selection and deployment under diverse constraints.

## 6. Conclusion

We introduce *BURMESE-SAN*, the first comprehensive benchmark for evaluating large language models on Burmese across NLU, NLR, and NLG tasks, constructed with high-quality, linguistically natural data spanning diverse domains.

Our evaluation reveals clear performance gaps between model families and generations, demonstrating that Burmese capability is strongly influenced by model architecture, instruction tuning, and training strategy rather than scale alone. In particular, Southeast Asian fine-tuned models - especially SEA-LION variants - consistently improve performance over generation while recent architectural advances such as MoE and reasoning-focused training further accelerate progress. Although commercial models currently achieve the highest overall scores, our results indicate that carefully tuned open-weight models can significantly narrow this gap, especially as Burmese-focused data and training strategies continue to improve.

Together, these findings underscore the importance of language-specific adaptation and position *BURMESE-SAN* as a robust foundation for future research, evaluation, and deployment of LLMs for Burmese and other low-resource languages.

## Acknowledgement

This research is supported by the National Research Foundation, Singapore, under its National Large Language Models Funding Initiative. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the National Research Foundation, Singapore.

The authors also would like to express their sincere gratitude to the native-speaker annotators and quality control contributors for their careful work and linguistic expertise. We also thank our internship students from King Mongkut’s University of Technology Thonburi (KMUTT), Htoo Myat Min Bo and Wira Ye Yint, for their valuable assistance with data annotation and quality assurance.

## Limitations

**Benchmarking with formal written Burmese prompt templates only** In this work, although the evaluation data are natural and conversational, we focus on evaluating models using formally written native Burmese prompt templates to ensure clarity, consistency, and grammatical correctness. Model performance on informal-style prompt templates may differ from what is observed in our benchmark. Extending *BURMESE-SAN* to include spoken and colloquial prompt templates could be a valuable direction for future work, enabling a more comprehensive evaluation of Burmese language models across different registers and contexts.

**Focus on Standard Burmese** We focus our study on the relatively better-resource standard Burmese from central region of Myanmar and do not include other notable dialects such as Arakanese (Rakhine) in the southwest, Tavoyan in the southeast, and Intha in the east, with others such as Yaw, Merguese (Myeik), and Palaw. Future work might explore translation or other community-driven data collection initiatives to extend coverage to dialects of Burmese language.

## Ethical Considerations

Our work on Burmese language technologies contributes to addressing key challenges in linguistic inclusion and improving technological accessibility for underrepresented language communities. The benchmark deliberately incorporates assessments of culturally specific knowledge to mitigate known biases in large language models. This emphasizes the importance of evaluating models within their cultural context rather than assuming universal applicability. Datasets included in *BURMESE-SAN*

are from publicly accessible sources and the authors check and include the License information of the datasets in the study.

For benchmark dataset quality assurance, a team of Burmese native speakers was involved in reviewing and annotating the data. The team, composed of university internship students, was recruited through faculty channels. Workload and compensation were communicated in advance and adhered to the university guidelines and regulatory requirements. Given the nature of the toxicity detection (TD) task, the authors and quality assurance team were exposed to potentially offensive material. Measures were taken to mitigate harm: annotators were encouraged to report inappropriate content and had the option to discontinue their work at any time during label verification and native revision.

We do not anticipate any negative social impacts arising from this study. The *BURMESE-SAN* dataset and accompanying codebase will be released under the Creative Commons Attribution Share-Alike 4.0 (CC-BY-SA 4.0) license.

## 7. Bibliographical References

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.

Sara Court and Micha Elsner. 2024. [Shortcomings of LLMs for low-resource translation: Retrieval and understanding are both the problem](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1332–1354, Miami, Florida, USA. Association for Computational Linguistics.

Jiaxin Duan, Fengyu Lu, and Junfei Liu. 2024. [Alleviating exposure bias in abstractive summarization via sequentially generating and revising](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 739–750, Torino, Italia. ELRA and ICCL.

Samuel Frontull and Thomas Ströhle. 2025. [Compensating for data with reasoning: Low-resource machine translation with llms](#).

Juraj Juraska, Daniel Deutsch, Mara Finkelstein, and Markus Freitag. 2024. [MetricX-24: The Google submission to the WMT 2024 metrics shared task](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 492–504, Miami, Florida, USA. Association for Computational Linguistics.

- Khin Aye. 2024. *Burmese Spoken Grammar*, 1st edition. Myanmar Language Commission.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Lei Li, Wei Liu, Marina Litvak, Natalia Vanetik, Jiacheng Pei, Yinan Liu, and Siya Qi. 2021. Subjective bias in abstractive summarization. *arXiv preprint arXiv:2106.10084*.
- Evan Miller. 2024. [Adding error bars to evals: A statistical approach to language model evaluations](#).
- Republic of the Union of Myanmar, Myanmar Language Commission. 2013. *Burmese Grammar*, 2nd edition. Myanmar Language Commission.
- Yewei Song, Lujun Li, Cedric Lothritz, Saad Ezzini, Lama Sleem, Niccolo Gentile, Radu State, Tegawendé F. Bissyandé, and Jacques Klein. 2025. [Is small language model the silver bullet to low-resource languages machine translation?](#)
- Open AI Team. 2025. [Openai gpt-5 system card](#).
- Gheorghe Comanici and et al. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#).
- DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Longxu Dou, Qian Liu, Fan Zhou, Changyu Chen, Zili Wang, Ziqi Jin, Zichen Liu, Tongyao Zhu, Cunxiao Du, Penghui Yang, Haonan Wang, Jiaheng Liu, Yongchi Zhao, Xiachong Feng, Xin Mao, Man Tsung Yeung, Kunat Pipatanakul, Fajri Koto, Min Si Thu, Hynek Kydlíček, Zeyi Liu, Qunshu Lin, Sittipong Sripaisarnmongkol, Kritaphad Sae-Khow, Nirattisai Thongchim, Taechawat Konkaew, Narong Borjindagoon, Anh Dao, Matichon Maneegard, Phakphum Arkaew, Zheng-Xin Yong, Quan Nguyen, Wannaphong Phatthiyaphaibun, Hoang H. Tran, Mike Zhang, Shiqi Chen, Tianyu Pang, Chao Du, Xinyi Wan, Wei Lu, and Min Lin. 2025. [Sailor2: Sailing in south-east asia with inclusive multilingual llms](#).

## Language Resource References

- Anthropic. 2025. [System card: Claude Opus 4 & Claude Sonnet 4](#). Accessed: 2026-02-18.
- Thura Aung, Ye Kyaw Thu, and Myat Noe Oo. 2024. [myocr: Optical character recognition for myanmar language with post-ocr error correction](#). In *2024 19th International Joint Symposium on Artificial Intelligence and Natural Language Processing (ISAI-NLP)*, pages 1–6.
- Shuai Bai and et al. 2025. [Qwen3-vl technical report](#).
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024. [The belebele benchmark: a parallel reading comprehension dataset in 122 language variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 749–775, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Ethnologue. 2019. Burmese. <https://web.archive.org/web/20190820164330/https://www.ethnologue.com/language/mya>. Accessed via Internet Archive on August 20, 2019.
- Ethnologue. 2024. Myanmar. <https://www.ethnologue.com/country/MM/>. Accessed: 27 Sept. 2024.
- Aaron Grattafiori and et al. 2024. [The llama 3 herd of models](#).
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. 2024. [OLMo: Accelerating the science of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.

- Jia Guo, Longxu Dou, Guangtao Zeng, Stanley Kok, Wei Lu, and Qian Liu. 2024a. [Sailcompass: Towards reproducible and robust evaluation for southeast asian languages](#).
- Jia Guo, Longxu Dou, Guangtao Zeng, Stanley Kok, Wei Lu, and Qian Liu. 2024b. [Sailcompass: Towards reproducible and robust evaluation for southeast asian languages](#). *arXiv preprint arXiv:2412.01186*.
- Zishan Guo, Renren Jin, Chuang Liu, Yufei Huang, Dan Shi, Linhao Yu, Yan Liu, Jiaxuan Li, Bojian Xiong, Deyi Xiong, et al. 2023. [Evaluating large language models: A comprehensive survey](#).
- Muhammad Usman Hadi, Rizwan Qureshi, Abbas Shah, Muhammad Irfan, Anas Zafar, Muhammad Bilal Shaikh, Naveed Akhtar, Jia Wu, Seyedali Mirjalili, et al. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XL-Sum: Large-scale Multilingual Abstractive Summarization for 44 Languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Zar Zar Hlaing, Ye Kyaw Thu, Thepchai Supnithi, and Ponrudee Netisopakul. 2022. Improving neural machine translation with pos-tag features for low-resource language pairs. *Heliyon*. Submitted Mar 5, 2022; revised Jun 25, 2022; accepted Aug 15, 2022; published online Aug 22, 2022.
- Aung Kyaw Htet and Mark Dras. 2025. [Myanmar xnli: Building a dataset and exploring low-resource approaches to natural language inference with myanmar](#).
- Xin Huang, Tarun Kumar Vangani, Minh Duc Pham, Xunlong Zou, Bin Wang, Zhengyuan Liu, and Ai Ti Aw. 2025. [Meralion-textllm: Cross-lingual understanding of large language models in chinese, indonesian, malay, and singlish](#).
- Mathias Jenny. 2021. [The national languages of msea: Burmese, thai, lao, khmer, vietnamese](#). In Paul Sidwell and Mathias Jenny, editors, *The Languages and Linguistics of Mainland Southeast Asia: A Comprehensive Guide*, pages 599–622. De Gruyter Mouton. Retrieved 2024-12-06.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Shengyi Jiang, Xiuwen Huang, Xiaonan Cai, and Nankai Lin. 2021. Pre-trained models and evaluation data for the myanmar language. In *The 28th International Conference on Neural Information Processing*, Cham. Springer International Publishing.
- Pride Kavumba, Naoya Inoue, Benjamin Heinzlerling, Keshav Singh, Paul Reiser, and Kentaro Inui. 2019. [When choosing plausible alternatives, clever hans can be clever](#). In *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, pages 33–42, Hong Kong, China. Association for Computational Linguistics.
- Nang Aeindray Kyaw, Ye Kyaw Thu, Thazin Myint Oo, Hutchatai Chanlekha, Manabu Okumura, and Thepchai Supnithi. 2024. [Enhancing hate speech classification in myanmar language through lexicon-based filtering](#). In *2024 21st International Joint Conference on Computer Science and Software Engineering (JCSSE)*, pages 316–323.
- Nathan Lambert and et al. 2025. [Tulu 3: Pushing frontiers in open language model post-training](#).
- Wei Qi Leong, Jian Gang Ngui, Yosephine Susanto, Hamsawardhini Rengarajan, Kengatharaiyer Sarveswaran, and William Chandra Tjhi. 2023. [Bhasa: A holistic southeast asian linguistic and cultural evaluation suite for large language models](#).
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re, Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekogul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2023. [Holistic evaluation of language models](#). *Transactions on Ma-*

- chine Learning Research*. Featured Certification, Expert Certification.
- Chaoqun Liu, Wenxuan Zhang, Jiahao Ying, Mahani Aljunied, Anh Tuan Luu, and Lidong Bing. 2025. [SeaExam and SeaBench: Benchmarking LLMs with local multilingual questions in Southeast Asia](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6119–6136, Albuquerque, New Mexico. Association for Computational Linguistics.
- Holy Lovenia, Rahmad Mahendra, Salsabil Maulana Akbar, Lester James V. Miranda, Jennifer Santoso, Elyanah Aco, Akhdan Fadhilah, Jonibek Mansurov, Joseph Marvin Imperial, Onno P. Kampman, Joel Ruben Antony Moniz, Muhammad Ravi Shulthan Habibi, Frederikus Hudi, Railey Montalan, Ryan Ignatius, Joanito Agili Lopo, William Nixon, Börje F. Karlsson, James Jaya, Ryandito Diandaru, Yuze Gao, Patrick Amadeus, Bin Wang, Jan Christian Blaise Cruz, Chenxi Whitehouse, Ivan Halim Parmonangan, Maria Khelli, Wenyu Zhang, Lucky Susanto, Reynard Adha Ryanda, Sonny Lazuardi Hermawan, Dan John Velasco, Muhammad Dehan Al Kautsar, Willy Fitra Hendria, Yasmin Moslem, Noah Flynn, Muhammad Farid Adilazuarda, Haochen Li, Johannes Lee, R. Damanhuri, Shuo Sun, Muhammad Reza Qorib, Amirbek Djanibekov, Wei Qi Leong, Quyet V. Do, Niklas Muennighoff, Tanrada Pansuwan, Ilham Firdausi Putra, Yan Xu, Tai Ngee Chia, Ayu Purwarianti, Sebastian Ruder, William Tjhi, Peerat Limkonchotiwat, Alham Fikri Aji, Sedrick Keh, Genta Indra Winata, Ruochen Zhang, Fajri Koto, Zheng-Xin Yong, and Samuel Cahyawijaya. 2024. [SEACrowd: A multilingual multimodal data hub and benchmark suite for Southeast Asian languages](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5155–5203, Miami, Florida, USA. Association for Computational Linguistics.
- Jann Railey Montalan, Jimson Paulo Layacan, David Demitri Africa, Richell Isaiah Flores, Michael T. Lopez II, Theresa Denise Magsajo, Anjanette Cayabyab, and William Chandra Tjhi. 2025. [Batayan: A filipino nlp benchmark for evaluating large language models](#).
- MULTICSD Project Team. 2025. [Burmese \(myanmar\)](#). <https://sites.google.com/view/multicsd/global-languages/burmese-myanmar>. Accessed: 2025-07-12.
- Raymond Ng, Thanh Ngan Nguyen, Huang Yuli, Tai Ngee Chia, Leong Wai Yi, Wei Qi Leong, Xianbin Yong, Jian Gang Ngui, Yosephine Susanto, Nicholas Cheng, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Adithya Venkatadri Hulagadri, Kok Wai Teng, Yeo Yeow Tong, Bryan Siow, Wei Yi Teo, Tan Choon Meng, Brandon Ong, Zhi Hao Ong, Jann Railey Montalan, Adwin Chan, Sajeban Antonyrex, Ren Lee, Esther Choa, David Ong Tat-Wee, Bing Jie Darius Liu, William Chandra Tjhi, Erik Cambria, and Leslie Teo. 2025. [SEA-LION: Southeast Asian languages in one network](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 512–526, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- John Okell. 2023. [Burmese \(Myanmar\): An Introduction To The Spoken Language, Book 1](#). CC0 1.0 Universal (Creative Commons Zero) License, Open Source.
- Team Olmo. 2025. [Olmo 3](#).
- Thazin Myint Oo, Thitipong Tanprasert, Ye Kyaw Thu, and Thepchai Supnithi. 2023. [Transfer and triangulation pivot translation approaches for burmese dialects](#). *IEEE Access*, 11:6150–6168.
- OpenAI. 2024. [Gpt-4 technical report](#).
- Libo Qin, Qiguang Chen, Yuhang Zhou, Zhi Chen, Yinghui Li, Lizi Liao, Min Li, Wanxiang Che, and Philip S. Yu. 2025. [A survey of multilingual large language models](#). *Patterns*, 6(1):101118.
- Republic of the Union of Myanmar, Ministry of Education. 2019. *University Technical Terms*. Department of Higher Education.
- Mya Ei San, Sasiporn Usanavasin, Ye Kyaw Thu, and Manabu Okumura. 2024. A study for enhancing low-resource thai-myanmar-english neural machine translation. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 23(4):1–24. Article No. 54. Submitted Aug 15, 2023; accepted with major revision Nov 4, 2023; accepted Jan 31, 2024; published online Apr 15, 2024.
- Guijin Son, Dongkeun Yoon, Juyoung Suk, Javier Aula-Blasco, Mano Aslan, Vu Trong Kim, Shayekh Bin Islam, Jaume Prats-Cristià, Lucía Tormo-Bañuelos, and Seungone Kim. 2024. [Mm-eval: A multilingual meta-evaluation benchmark for llm-as-a-judge and reward models](#).
- Yosephine Susanto, Adithya Venkatadri Hulagadri, Jann Railey Montalan, Jian Gang Ngui, Xian Bin Yong, Weiqi Leong, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Yifan Mai, and

- William Chandra Tjhi. 2025. [Sea-helm: South-east asian holistic evaluation of language models](#).
- DeepSeek-AI Team. 2025a. [Deepseek-v3 technical report](#).
- Gemini Team. 2025b. [Gemini: A family of highly capable multimodal models](#).
- Gemma Team. 2025c. [Gemma 3 technical report](#).
- NLLB Team. 2022. No Language Left Behind: Scaling Human-Centered Machine Translation.
- NLLB Team. 2024a. [Scaling neural machine translation to 200 languages](#). *Nature*, 630(8018):841–846.
- OpenAI Team. 2024b. [Gpt-4o system card](#).
- OpenAI Team. 2025d. [gpt-oss-120b gpt-oss-20b model card](#).
- Qwen Team. 2025e. [Qwen2.5 technical report](#).
- Qwen Team. 2025f. [Qwq-32b: Embracing the power of reinforcement learning](#).
- Ye Kyaw Thu, Thura Aung, and Thepchai Supnithi. 2023. Neural sequence labeling based sentence segmentation for myanmar language. In *The 12th Conference on Information Technology and Its Applications*, pages 285–296, Cham. Springer Nature Switzerland.
- Ashmal Vayani, Dinura Dissanayake, Hasindri Watawana, Noor Ahsan, Nevasini Sasikumar, Omkar Thawakar, Henok Biadgign Ademtew, Yahya Hmaiti, Amandeep Kumar, Kartik Kuckreja, Mykola Maslych, Wafa Al Ghallabi, Mihail Mihaylov, Chao Qin, Abdelrahman M Shaker, Mike Zhang, Mahardika Krisna Ihsani, Amiel Esplana, Monil Gokani, Shachar Mirkin, Harsh Singh, Ashay Srivastava, Andre Hamerlik, Fathinah Asma Izzati, Fadillah Adamsyah Maani, Sebastian Cavada, Jenny Chim, Rohit Gupta, Sanjay Manjunath, Kamila Zhumakhanova, Feno Heriniaina Rabevohitra, Azril Amirudin, Muhammad Ridzuan, Daniya Kareem, Ketan More, Kunyang Li, Pramesh Shakya, Muhammad Saad, Amirpouya Ghasemaghaei, Amirbek Djanibekov, Dilshod Azizov, Branislava Jankovic, Naman Bhatia, Alvaro Cabrera, Johan Obando-Ceron, Olympiah Otieno, Fabian Farestam, Muztoba Rabbani, Sanoojan Baliah, Santosh Sanjeev, Abduragim Shtanchaev, Maheen Fatima, Thao Nguyen, Amrin Kareem, Toluwani Aremu, Nathan Xavier, Amit Bhatkal, Hawau Toyin, Aman Chadha, Hisham Cholakkal, Rao Muhammad Anwer, Michael Felsberg, Jorma Laaksonen, Tamar Solorio, Monojit Choudhury, Ivan Laptev, Mubarak Shah, Salman Khan, and Fahad Khan. 2025. [All languages matter: Evaluating llms on culturally diverse 100 languages](#).
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy Chen. 2024. [SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 370–390, Mexico City, Mexico. Association for Computational Linguistics.
- Bryan Wilie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, and Ayu Purwarianti. 2020. [IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 843–857, Suzhou, China. Association for Computational Linguistics.
- An Yang and et al. 2025. [Qwen3 technical report](#).
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#).
- Wint Theingi Zaw, Ye Kyaw Thu, Zar Zar Hlaing, and Thepchai Supnithi. 2022. English to burmese machine translation with asian pivot languages. *Journal of Intelligent Informatics and Smart Technology*, Oct 2nd Issue:26–1–26–9. Submitted Feb 21, 2022; accepted Aug 30, 2022; published online Oct 31, 2022.
- Yiran Zhao, Chaoqun Liu, Yue Deng, Jiahao Ying, Mahani Aljunied, Zhaodonghui Li, Lidong Bing, Hou Pong Chan, Yu Rong, Deli Zhao, and Wenxuan Zhang. 2025. [Babel: Open multilingual large](#)

language models serving over 90% of global speakers. *arXiv preprint arXiv:2503.00865*.

## A. Overview of Tasks and Datasets

As described in Section 3, we selected seven tasks for inclusion in *BURMESE-SAN*. Table 1 lists the source datasets, their adaptations, and usage, all consistent with the original intent.

- **Abstractive Summarization (AS)** In this task, an LLM is given a paragraph and must generate a concise sentence summarizing its content. The evaluation focuses not only on identifying the key information but also on paraphrasing the content coherently. We use XL-Sum (Hasan et al., 2021) for this task, which contains annotated article-summary pairs.
- **Causal Reasoning (CR)** This task requires the LLM to identify the causal relationship between events. Given a premise and a set of statements, the model must determine which statement represents the cause or effect of the premise. We translate from scratch and employ Balanced COPA (Kavumba et al., 2019), designed to evaluate commonsense causal reasoning with paired alternatives.
- **Machine Translation (MT)** Here, an LLM is provided with text in one language and is expected to translate it into another language. In this work, we evaluate both *English to Burmese* and *Burmese to English* translations using the FLORES+ dataset (Team, 2024a), which includes translations across multiple languages and domains.
- **Natural Language Inference (NLI)** This classification task requires the LLM to determine the relationship between two sentences (X and Y) as one of the following: (a) X implies Y, (b) X contradicts Y, or (c) X neither implies nor contradicts Y. We use myXNLI (Htet and Dras, 2025), which provides human-annotated examples for cross-lingual inference evaluation.
- **Question Answering (QA)** In this task, an LLM is given a passage and a question and must select the span from the passage that answers the question. We use Belebele (Bandarkar et al., 2024), a multiple-choice reading comprehension dataset designed to evaluate passage understanding.
- **Toxicity Detection (TD) and Sentiment Analysis (SA)** Both tasks involve analyzing natural language text. TD requires detecting hate speech or abusive language, while SA involves classifying sentiment as positive, negative, or neutral. We use myHateSpeech (Kyaw et al., 2024) for TD and GKLMIP-mya (Jiang et al., 2021) for SA.

Task	Dataset	License
QA	Belebele (Bandarkar et al., 2024)	CC BY-NC 4.0
SA	GKLMIP-mya (Jiang et al., 2021)	Unknown
TD	myHateSpeech (Kyaw et al., 2024)	CC BY-NC-SA 4.0
CR	Balanced COPA (Kavumba et al., 2019)	CC BY 4.0
NLI	myXNLI (Htet and Dras, 2025)	CC BY-NC 4.0
AS	XL-Sum (Hasan et al., 2021)	CC BY-NC-SA 4.0
MT	FLORES+ (Team, 2024a)	CC BY-SA 4.0

Table 6: Licenses for all datasets included in *BURMESE-SAN*, with corresponding references.

## B. Dataset Quality Assurance

To ensure dataset quality, we employed bilingual native speakers (ages 18–25), primarily university students enrolled in international programs in Thailand.

- (ENG) Assist doctors while providing patient care.
- (Acceptable: Two sentences) ဆရာဝန်တွေကို ကူညီပါ။ (Help doctors) လူနာတွေကိုလည်း တစ်ချိန်တည်းမှာ ကြည့်ရှုပေးပါ။ (Help patients at the same time)
- (Acceptable: One sentence different order with ENG) လူနာတွေကို ကြည့်ရှုရင်း ဆရာဝန်တွေကို ကူညီပါ။ (While helping patients, help doctors)
- (Acceptable: Same order with ENG, incorrect grammar in Burmese, same context) ဆရာဝန်တွေကို ကူညီပါ လူနာတွေကို ကြည့်ပေးရင်းနဲ့
- (Not acceptable: wrong context but correct grammar in Burmese) ဆရာဝန်တွေကို ကြည့်ပေးရင်းနဲ့ (While caring doctors) လူနာတွေကို ကူညီပါ (Help patients)

Figure 3: Acceptable and Not Acceptable Grammar Errors in the Dataset.

- Consonant error: These errors occur while misusing consonants, vowels, and independent vowels. (e.g ဥသြ → ဥသြ, သြဂုတ်လ → သြဂုတ်လ, ကြက်ဥ → ကြက်ဥ)
- Dialect error: These errors occur due to language variety based on community or particular area (e.g စပ်စုလတ်သာအေဖယ် → စပ်စုလိုက်သာအေဖယ်)
- Encoding error: Sometimes Burmese data comes with Zawgyi (Non-Unicode) encoding: (e.g စပ်စုလတ်သာအေဖယ် → စပ်စုလိုက်သာအေဖယ်)
- Phonetic error: Pronunciation of the misspelled word is the same with the intended correct word. (e.g ခေါင်းညိမ့် → ခေါင်းညိတ်, မဟုတ်ပဲ → မဟုတ်ဘဲ)
- Typographic error: (e.g နိုင် → နိုင်)
- Sequence error: Typed in incorrect order. (e.g ကေမြ → ကြေမြ)
- Short error: Just short form in daily messaging (e.g အင် → အခြောက်)
- Slang error: Some words are misspelled and migrated as slangs (e.g မိဉ်း မြို့)
- Stack error: Some adopted words from Pali are stack words and can be misspelled based on the above errors (e.g မင်လာ → မင်လာ)

Figure 4: Different Types of Spelling Errors.

Each task may include different types of linguistic issues. QC members need to fix Grammar and Spelling errors, and the fixed datasets are later used for evaluating LLMs. Figure 3 and 4 shows acceptable and unacceptable grammar errors and spelling errors examples. For the grammatical references, we used official guideline published by [Republic of the Union of Myanmar, Myanmar Language Commission 2013](#) and spoken grammar written by [Khin Aye 2024](#).

To ensure the quality and reliability of the translated datasets, **three** QC members evaluated the English–Burmese text pairs for the QA, MT, AS, and CR tasks. Each member assigned a score of 0 (Disagree) or 1 (Agree) for each evaluation criterion, which varied by task. The definition of criteria are as shown in the Table 7.

In addition to dataset quality checks, QC members responsible for translation were instructed to focus on the following issues:

Criteria	Tasks	Definition
Completeness	All	Is all the intended information from the source instruction retained in the translated instruction?
Fluency	All	Is the output grammatically correct and natural in Burmese?
Sensibility	All	Is the translation logical and sensible based on the context of the original instruction?
Faithfulness	Translation	Does the translated instruction stay true to the meaning of the original instruction?
Relevance	Summarization	Does the output (summary) contain only the most important content?
Coherence	Summarization	Is the output (summary) logically structured and easy to follow given the context of the original instruction?

Table 7: Quality Evaluation Criteria for *BURMESE-SAN*

Task	Criterion	Joint Agreement (%)
QA	Completeness	100.00
	Fluency	97.50
	Sensibility	93.33
MT	Completeness	100.00
	Fluency	99.76
	Sensibility	99.88
	Grammaticality	99.76
	Faithfulness	99.88
AS	Completeness	100.00
	Fluency	100.00
	Sensibility	100.00
	Relevance of Summary	100.00
	Fluency of Summary	100.00
	Coherence of Summary	100.00
CR	Completeness	100.00
	Fluency	97.00
	Sensibility	97.75

Table 8: Joint agreement (%) for each evaluation criterion across tasks requiring (re-)translation. Scores are before revision by native speakers.

- **Literal Translation:** Avoid overly direct word-for-word rendering that neglects the target language’s natural usage and style.
- **Cultural Mismatch:** Identify translations that sound unnatural due to irrelevant or inappropriate cultural references.
- **Incomplete Translation:** Ensure that no parts of the source text are omitted or skipped.
- **Misinterpretation:** Verify that the intended meaning of the source text is accurately preserved in the translation.

After evaluation, we revised all translated datasets by reviewing samples that did not receive full agreement among the three annotators. As shown in Table 8, most criteria achieved high agreement scores - many reaching 100% - while a few, particularly for Fluency and Sensibility in QA and CR tasks, were slightly lower. These scores reflect the results after re-translation but before manual revision. To ensure high quality and consistency, samples without full agreement were further revised by native speakers. The final dataset contains only samples with full consensus across all evaluation criteria. Dataset statistics, including class distribution are provided in Table 2.

## C. Challenges with Developing a Burmese Benchmark

### C.1. Issues with Native-Sourced and Native-Translated Data

**Inconsistency in Technical Terms** A common challenge we encountered in Burmese datasets, especially those created or translated by native speakers, is the inconsistent use of loanwords and technical terms. While a Burmese dictionary for scientific and technical vocabulary exists ([Republic of the Union of Myanmar, Ministry of Education, 2019](#)), it is not widely adopted, and many modern terms are missing. As shown in Figure 5 (example a), different translators use different styles. For example, MYA1 and MYA2 provide different translations for the word “Theoretical” and use different transliterations for the city name “Beijing.”

(a) ENG: I am going to visit **Beijing** this summer for studying **Theoretical Physics**

MYA1: ဒီနွေရာသီမှာ သီအိုရီ ရူပဗေဒကို လေ့လာဖို့ ဘေဂျင်းကို သွားမယ်။

MYA2: ဒီနွေရာသီမှာ သဘောတရားရေး ရူပဗေဒကို လေ့လာဖို့ ဖေဂျင်းကို သွားမယ်။

MYA3: ဒီနွေရာသီမှာ ဘေဂျင်းကို သွားမယ် သီအိုရီ ရူပဗေဒကို လေ့လာဖို့။

(b) ENG: Computer Science is a branch of Science.

MYA1: ကွန်ပျူတာသိပ္ပံသည် သိပ္ပံပညာ၏ ဘာသာရပ်ခွဲတစ်ခုဖြစ်သည်။

MYA2: ကွန်ပျူတာသိပ္ပံဟာ သိပ္ပံရဲ့ ဘာသာရပ်ခွဲတစ်ခု ဖြစ်ပါတယ်။

MYA3: ကွန်ပျူတာသိပ္ပံဟာ သိပ္ပံပညာရဲ့ ဘာသာရပ်ခွဲတစ်ခုပါ။

Figure 5: Examples of variation in Burmese translations by native speakers. (a) Differences in technical term usage, transliteration, and word order. (b) Differences in particle choice and formality.

**Inconsistency in Syntax** We encounter another problem, which is syntax inconsistency. As shown in 5(a), MYA3’s translation follows a similar structure to MYA1 but uses a different word order. Although the syntax is not strictly correct, the meaning remains understandable. MYA3 places the phrase “I am going to visit Beijing” before “for studying Theoretical Physics” to emphasize the trip itself. In Figure 5 (example b), the translations also vary in particle usage. Burmese has multiple particles with similar meanings, and different translators make different stylistic choices. For instance, MYA2 and MYA3 use colloquial language, whereas MYA1 employs formal expressions.

These examples illustrate the broader challenges of working with native-sourced and native-translated Burmese data. The lack of commonly followed standards leads to inconsistencies in terminology. For the syntax, although there are writing standards for both formal and informal styles, the standards are rarely followed. Burmese is a rich and flexible language, and in practice, many speakers do not strictly adhere to formal grammar or standardized syntax, especially in informal contexts like social media. This makes it difficult to ensure consistency, even among native speakers. Such variation can impact the quality and reliability of datasets, particularly for downstream tasks like machine translation or causal reasoning. To address this, we conducted a cross-revision process by native speakers reviewing and refining each others’ annotations, or translations to improve consistency, clarity, and alignment with the original meaning from original native-sourced or original native-translated datasets. With this approach, we ensure the high-quality and consistency of our datasets, which reflect the real-world application.

### C.2. Issues with English-sourced and English-adapted Data

English-sourced and English-adapted datasets were initially translated using either semi-automatic methods or manual translation. However, we found that automatic translations, such as those from Google Translate, while grammatically correct, were often unnatural, disfluent, and showed signs of translation. Many previous works also demonstrate the failure in low-resource machine translations, which results in low-quality ([Frontull and Ströhle, 2025](#); [Court and Elsner, 2024](#)) or a lack of localization ([Song et al., 2025](#)). To address this, we applied retranslation to ensure the samples sounded more natural and native.

For the English-sourced dataset, Balanced COPA (Kavumba et al., 2019), there is no Burmese translation in the previous works. Therefore, we need to translate the dataset to Burmese. We used Google Translate, but several issues arose due to literal translation of idioms, incorrect word choices, and reversed meanings. These issues impacted the semantic fidelity of the dataset, especially in the Balanced COPA translation sourced from English. Therefore, we translated the corpus with a bilingual native speaker manually and rated the translation quality, then revised the translation errors again following the procedures discussed in Section 3.3 and Appendix B. These problems highlight the need for culturally and contextually aware translations.

Similar issues were found in the original translated version of English-adapted datasets like Flores+ and Belebele, where mistranslations, with unclear question intents and word-for-word translations, often lead to loss of meaning, especially in culturally sensitive or nuanced cases. In another English-adapted dataset, XL-Sum, in the original translation, we identified several quality issues, including inaccurate summaries, missing key information, inconsistent or misleading titles, and incomplete articles. These problems may affect the factual reliability and coherence of the data, highlighting the need for thorough human validation.

Despite re-translation efforts, some samples still did not fully meet fluency and sensibility criteria, as shown in Table 8. To ensure high-quality and consistent samples for evaluation, we conducted additional manual revisions. In *BURMESE-SAN*, we addressed these challenges through a rigorous translation and revision process in collaboration with native speakers.

## D. Dataset Examples Prompt Templates, and Models Evaluated

Task	Example
QA	<p><b>Text:</b> ဥက္ကာခဲများသည် ပရိုတင်း နှင့် ရှင်သန်ရေး အထောက်အကူပြုများ ဖြစ်စေနိုင်သည် အော်ဂဲနစ်ပစ္စည်းနှင့်အတူ ကမ္ဘာသို့ ရေအရင်းမြစ်ကို သယ်ဆောင်ပေးခြင်း ဖြစ်နိုင်သည်။ လွန်ခဲ့သော နှစ်များစွာ အကြာ ကမ္ဘာကြီးနှင့် ကြယ်တံခွန်များ တိုက်မိခဲ့စဉ် အချိန် ကတည်းက သိပ္ပံပညာရှင်များသည် ဂြိုဟ်များဖြစ်ပေါ်လာပုံ၊ အထူးသဖြင့် ကမ္ဘာကြီး ဖြစ်ပေါ်လာပုံကို နားလည်လို ကြပါသည်။</p> <p><b>Question:</b> သိပ္ပံပညာရှင်များ ရှာဖွေသိရှိလိုသည့် တစ်စုံတစ်ရာမှာ အဘယ်နည်း။</p> <p><b>Choices:</b>            (က) ကမ္ဘာကြီးနှင့် ကြယ်တံခွန်များ အပြင်းအထန်တိုက်မိခဲ့ချိန် (ခ) ပရိုတိန်းများ ဖြစ်ပေါ်ပုံ            (ဂ) သဘာဝဒြပ်အကြောင်း (ဃ) ကမ္ဘာကြီး ဖြစ်တည်လာပုံ</p> <p><b>Answer:</b> ဃ</p>
SA	<p><b>Text:</b> ပစ္စည်းကလက်ဆောင်ရောပါလို့ကောင်းပါတယ်။</p> <p><b>Label:</b> ကြားနေခံစားချက် (neutral)</p>
TD	<p><b>Text:</b> ညနေ သွား ကျိတ် မယ် ဗျာ လုပ် လိုက် တော့ 😊</p> <p><b>Label:</b> သန့်ရှင်း (Clean)</p>
CR	<p><b>Text:</b> ကျွန်ုပ်၏ကွန်ပျူတာအသံသည် အလုပ်မလုပ်ပါ။</p> <p><b>Question:</b> အကျိုး (effect)</p> <p><b>Choices:</b> (0) စပီကာအသံတွေ တပ်ဆင်ခဲ့တယ်။ (1) ကျွန်ုပ်၏ ဒေတာအားလုံးဆုံးရှုံးသွားခဲ့သည်။</p> <p><b>Label:</b> (0)</p>
NLI	<p><b>Sentence 1:</b> ဤစာကို လက်ခံရရှိသည့် လူတိုင်း ၁၈ ဒေါ်လာလောက်သာ ပေးခဲ့လျှင်။</p> <p><b>Sentence 2:</b> ဤစာကိုလက်ခံရရှိသူတိုင်း- သင်တို့၏ ပိုက်ဆံကို မလှူဒါန်းပါနဲ့၊ အဲဒါ လိမ်လည်မှုတစ်ခုဖြစ်တယ်။</p> <p><b>Label:</b> ဆန့်ကျင် (contradiction)</p>
AS	<p><b>Article:</b> တကယ့်ကို သမိုင်းဝင်တဲ့ အခိုက်အတံ့ ဖြစ်ကြောင်း ဂျွန်ကယ်ရီ ပြော။ အမေရိကန် နိုင်ငံခြားရေး ဝန်ကြီး တစ်ဦး အနေနဲ့ နှစ်ပေါင်း ၇၀ အတွင်း ကျူးဘားကို ပထမဆုံး သွားရောက်သူ ဖြစ်လာတဲ့ ဂျွန်ကယ်ရီက ဟာဗားနား မြို့က အမေရိကန် သံရုံး ဖွင့်ပွဲ အလံတင် အခမ်းအနားကို ကြီးကြပ် ခဲ့ပါတယ်။ ၁၉၆၁ ခုနှစ်က အလံ ဖြုတ်ချခဲ့တဲ့ မရိန်းတပ်သား ၃ ဦးကပဲ အခု သံရုံး ဖွင့်ပွဲမှာ အလံတင် ခဲ့ပါတယ်။ အခုလို အလံလွှင့်တင်မှုဟာ သမိုင်းမှာ အရေးပါတဲ့ အခိုက်အတံ့ ဖြစ်တယ်လို့ မစ္စတာ ကယ်ရီက ဒီကနေ့ သောကြာနေ့ အခမ်းအနားမှာ ပြောခဲ့ပါတယ်။ ဒါပေမယ့် ကျူးဘားမှာ နိုင်ငံရေး အပြောင်းအလဲ ဖြစ်ဖို့ ဖိအားပေးမှုတွေကိုတော့ အမေရိကန်ဘက်က ရပ်တန့်မှာ မဟုတ်ဘူးလို့ သူက သတိပေးခဲ့ပါတယ်။ အမေရိကန်ဘက်က ကျူးဘား သံရုံးကိုတော့ ဝါရှင်တန်မြို့မှာ ပြီးခဲ့တဲ့ လက ဖွင့်ခဲ့ပြီး ဖြစ်ပါတယ်။ ဒါပေမယ့် အမေရိကန်ဘက်က ကုန်သွယ်ရေး ပိတ်ဆို့ထားမှုကို မပယ်ဖျက် သေးတဲ့ ကိစ္စနဲ့ ပတ်သက်ပြီး ကျူးဘား သမ္မတ ဖီဒယ် ကက်စထရိုက လူသိရှင်ကြား ပြောခဲ့ပါတယ်။</p> <p><b>Summary:</b> ကျူးဘား နိုင်ငံမှာ ၅၄ နှစ်ကျော်ကြာ ပိတ်ထားခဲ့တဲ့ အမေရိကန် သံရုံးကို ဒီကနေ့ သောကြာနေ့မှာ ပြန်လည် ဖွင့်လှစ် လိုက်ပါတယ်။</p>
MT	<p><b>ENG:</b> He recently lost against Raonic in the Brisbane Open.</p> <p><b>MYA:</b> သူသည် မကြာသေးမီက ဘရစ်ဘိတ်အိုးပင်းပြိုင်ပွဲတွင် ရောအိုနစ်ကို ရှုံးနိမ့်ခဲ့ရသည်။</p>

Figure 6: Example Data Samples for each task in *BURMESE-SAN*.

သင့်ကို စာပိုဒ်တစ်ပိုဒ်၊ မေးခွန်းတစ်ခု နှင့် ရွေးချယ်စရာ အဖြေ ဝေဖန် ပေးထားပါလိမ့်မည်။ စာပိုဒ်ကိုအခြေခံပြီး ပေးထား သော ရွေးချယ်စရာများထဲမှ တစ်ခုကို ရွေးချယ်၍ ဖြေဆိုပါ။

A paragraph to you You will be given a question and four possible answers. Choose one of the given options and answer based on the passage.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အဖြေ- \$OPTION  
\$OPTION နေရာတွင် သင်ရွေးချယ်ထားသော အဖြေကို အစားထိုးထည့်ပါ။ အဖြေအတွက် က၊ ခ၊ ဂ သို့မဟုတ် ဃ အက္ခရာကို အသုံးပြုပါ။(fewshot\_examples)

Answer using only the form below:  
Answer: \$OPTION  
Substitute your chosen answer in place of \$OPTION. For the answer, b. Use the letter c or d. (fewshot\_examples)

စာပိုဒ်-  
...  
{text}  
...  
မေးခွန်း- {question}  
က- {choice1}  
ခ- {choice2}  
ဂ- {choice3}  
ဃ- {choice4}

Paragraph:  
...  
{text}  
...  
Question: {question}  
A: {choice1}  
B: {choice2}  
C: {choice3}  
D: {choice4}

အောက်ပါဝါကျစာကြောင်း၏ စိတ်ဓာတ်ချက် မှာ အဘယ်နည်း။ စိတ်ဓာတ်ချက် ဆိုသည်မှာ ရေးသား ပြောဆိုစေ့စပ် ထိုသူ၏ စိတ် ထဲတွင် ဖြစ်နေသော အပြုသဘော၊ အယူသဘော၊ သို့မဟုတ် အမြဲမဟုတ် အယူကွဲကွဲသော ကြားနေ ဓာတ်ချက် ဖြစ်သည်။ ပေးထားသော စာကြောင်းသည် အပြုသဘော၊ အယူကွဲကွဲသော သို့မဟုတ် ကြားနေ မည်သည့် သဘောရှိသည်ကို စကားလုံးတစ်လုံး တည်းဖြင့်သာ အမျိုးအစား သတ်မှတ်ပါ။

What is the sentiment of the following sentence? Sentiment refers to the positive, negative, or neutral feelings that the person has while writing or speaking. Determine the type of sentiment in the given sentence using just one word: positive, negative, or neutral.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အဖြေ- \$OPTION  
\$OPTION နေရာတွင် သင်ရွေးချယ်ထားသော အဖြေကို အစားထိုး ထည့်ပါ။ (fewshot\_examples)

ဝါကျစာကြောင်း-  
...  
{text}  
...  
Sentence:  
...  
{text}  
...  
Use the following format to respond: Answer: \$OPTION Replace \$OPTION with your chosen answer.(fewshot\_examples)

(a) Question Answering

(b) Sentiment Analysis

သင်သည် အဆိပ်အတောက် (toxicity) ရှာဖွေစစ်ဆေးသည့်စနစ် တစ်ခုဖြစ်သည်။ အညွှန်းများကို အောက်ပါအတိုင်း အဓိပ္ပာယ်ဖွင့် ဆိုထားပါသည်။

You are a toxicity detection system. The labels are defined as follows:  
Clean: No toxicity at all.  
Hate speech: includes toxicity or attacks directly targeting someone.

သန့်ရှင်း- အဆိပ်အတောက် (toxicity) လုံးဝမပါဝင်ပါ။  
အမုန်းကား- အဆိပ်အတောက် (toxicity) သို့မဟုတ် တစ်စုံတစ်ခုလုံးအား တိုက်ခိုက် ပစ်မှတ်ထား တိုက်ခိုက်ခြင်းများ ပါဝင်သည်။

According to the above index definitions, Classify the sentences below as clean or hate speech using only one word.

အထက်ပါ အညွှန်း အဓိပ္ပာယ်ဖွင့်ဆိုချက်များအရ အောက်ဖော်ပြ ပါ ဝါကျစာကြောင်း ကို သန့်ရှင်း သို့မဟုတ် အမုန်းကား ဟူ၍ စကားလုံးတစ်လုံးတည်းဖြင့်သာ အမျိုးအစား သတ်မှတ်ပါ။

Answer using only the form below:  
Answer: LABEL  
Replace LABEL with the selected label.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အဖြေ- \$LABEL  
\$LABEL ကို ရွေးချယ်ထားသည့် အညွှန်းဖြင့် အစားထိုးပါ။  
ဝါကျစာကြောင်း- {text}

sentence string: {text}

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အဖြေ- \$OPTION  
\$OPTION နေရာတွင် သင်ရွေးချယ်ထားသော အဖြေကို အစားထိုးထည့်ပါ။ အဖြေအတွက် က သို့မဟုတ် ခ အက္ခရာကို အသုံးပြုပါ။ (fewshot\_examples)

Answer using only the form below:  
Answer: \$OPTION  
Substitute your chosen answer in place of \$OPTION. Use the letter a or b for the answer.(fewshot\_examples)

ပေးထားသော အဖြေအနက် အခြေခံ၍ အောက်ပါရွေးချယ်စရာ များထဲမှ မည်သည့်အရာက (question\_translated) ဖြစ်နိုင်ခြေ ပို၍ များသနည်း။  
အဖြေအနေ-  
...  
{text}  
...  
အောက်ပါရွေးချယ်စရာများမှ အကောင်းဆုံးအဖြေကို ရွေးချယ် ပါ-  
က- {choice1}  
ခ- {choice2}

Based on the given situation, which of the following options is (question\_translated) most likely?  
Condition:  
...  
{text}  
...  
Choose the best answer from the following options:  
A: {choice1}  
B: {choice2}

(c) Toxicity Detection

(d) Causal Reasoning

သင့်ကို SENTENCE\_1 နှင့် SENTENCE\_2 ဟူသော ဝါကျစာကြောင်း နှစ် ခုကို ပေးထားပါမည်။ SENTENCE\_1 နှင့် SENTENCE\_2 တို့အတွက် အောက်ဖော်ပြချက်များထဲမှ မည်သည့်အချက်က အကောင်းဆုံး ကိုက်ညီမှုရှိသည်ကို ဆုံးဖြတ်ပါ။

You will be given two sentences called SENTENCE\_1 and SENTENCE\_2. Decide which of the following statements best matches SENTENCE\_1 and SENTENCE\_2.  
A: If SENTENCE\_1 is true, then SENTENCE\_2 must be true.  
B: SENTENCE\_1 contradicts SENTENCE\_2.  
C: When SENTENCE\_1 is correct, SENTENCE\_2 may or may not be valid.

က- SENTENCE\_1 မှန်ကန်သည့် SENTENCE\_2 သည် မုချမှန်ကန်ရ မည်။  
ခ- SENTENCE\_1 သည် SENTENCE\_2 ကို ဆန့်ကျင်သည်။  
ဂ- SENTENCE\_1 မှန်ကန်သောအခါ၊ SENTENCE\_2 သည် မှန်ကန်နိုင် သလို မမှန်ကန်သလည်း ရှိနိုင်ပါသည်။

Answer using only the form below:  
Answer: \$OPTION  
Substitute your chosen answer in place of \$OPTION. For the answer, Use the letter b or c.(fewshot\_examples)

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အဖြေ- \$OPTION  
\$OPTION နေရာတွင် သင်ရွေးချယ်ထားသော အဖြေကို အစားထိုးထည့် ပါ။ အဖြေအတွက် က၊ ခ သို့မဟုတ် ဂ အက္ခရာကို အသုံးပြုပါ။ (fewshot\_examples)

SENTENCE\_1-  
...  
{sentence1}  
...  
SENTENCE\_2-  
...  
{sentence2}

SENTENCE\_1:  
...  
{sentence1}  
...  
SENTENCE\_2:  
...  
{sentence2}

အောက်ပါ မြန်မာဆောင်းပါးကို ဝါကျ စာကြောင်း တစ် ကြောင်း သို့မဟုတ် နှစ်ကြောင်းပါသော စာပိုဒ်တစ်ပိုဒ်အဖြစ် အကျဉ်းချုပ်ဖော်ပြပါ။

Summarize the following Burmese article as a paragraph of one or two sentences.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
အကျဉ်းချုပ်- \$SUMMARY  
\$SUMMARY ကို အကျဉ်းချုပ်ဖြင့် အစားထိုးပါ။

Answer using only the form below:  
Summary: \$SUMMARY  
Replace \$SUMMARY with summary.

Article:  
...  
{text}  
...

ဆောင်းပါး-  
...  
{text}  
...

(e) Natural Language Inference

(f) Abstractive Summarization

အောက်ပါစာသားကို အင်္ဂလိပ်ဘာသာသို့ ဘာသာပြန်ပါ။

Translate the following text into English.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
ဘာသာပြန်ချက်- \$TRANSLATION  
\$TRANSLATION ကို ဘာသာပြန်ထားသော စာသားဖြင့် အစားထိုးပါ။

Answer using only the form below:  
Translation: \$TRANSLATION  
Replace \$TRANSLATION with translated text.

Text:  
...  
{text}  
...

စာသား-  
...  
{text}  
...

အောက်ပါစာသားကို မြန်မာဘာသာသို့ ဘာသာပြန်ပေးပါ။

Please translate the following text into Burmese.

အောက်ဖော်ပြပါ ပုံစံကိုသာ အသုံးပြု၍ ဖြေဆိုပါ-  
ဘာသာပြန်ချက်- \$TRANSLATION  
\$TRANSLATION ကို ဘာသာပြန်ထားသော စာသားဖြင့် အစားထိုးပါ။

Answer using only the form below:  
Translation: \$TRANSLATION  
Replace \$TRANSLATION with translated text.

Text:  
...  
{text}  
...

စာသား-  
...  
{text}  
...

(g) Machine Translation (to English)

(h) Machine Translation (to Myanmar)

Figure 7: Prompt templates used for BURMESE-SAN. English prompt versions are also provided.

Models	Source Link	No. params
<i>Instruction Tuning Models</i>		
ERNIE 4.5	<a href="https://huggingface.co/baidu/ERNIE-4.5-300B-A47B-PT">https://huggingface.co/baidu/ERNIE-4.5-300B-A47B-PT</a>	300B MoE
Qwen 3 A22B	<a href="https://huggingface.co/Qwen/Qwen3-235B-A22B-Instruct-2507">https://huggingface.co/Qwen/Qwen3-235B-A22B-Instruct-2507</a>	235B MoE
Llama 4 Maverick	<a href="https://huggingface.co/meta-llama/Llama-4-Maverick-17B-128E-Instruct">https://huggingface.co/meta-llama/Llama-4-Maverick-17B-128E-Instruct</a>	400B MoE
DeepSeek V3.1	<a href="https://huggingface.co/deepseek-ai/DeepSeek-V3.1">https://huggingface.co/deepseek-ai/DeepSeek-V3.1</a>	671B MoE
SEA-LION v4 (Qwen 3)	<a href="https://huggingface.co/aisingapore/Qwen-SEA-LION-v4-32B-IT">https://huggingface.co/aisingapore/Qwen-SEA-LION-v4-32B-IT</a>	32B
Gemma 3 VL	<a href="https://huggingface.co/google/gemma-3-27b-it">https://huggingface.co/google/gemma-3-27b-it</a>	27B
SEA-LION v4 (Gemma 3)	<a href="https://huggingface.co/aisingapore/Gemma-SEA-LION-v4-27B-IT">https://huggingface.co/aisingapore/Gemma-SEA-LION-v4-27B-IT</a>	27B
Llama 4 Scout	<a href="https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E-Instruct">https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E-Instruct</a>	109B MoE
Qwen 3 Next	<a href="https://huggingface.co/Qwen/Qwen3-Next-80B-A3B-Instruct">https://huggingface.co/Qwen/Qwen3-Next-80B-A3B-Instruct</a>	80B MoE
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-32B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-32B-Instruct</a>	32B
Kimi K2 Instruct 0905	<a href="https://huggingface.co/moonshot-ai/Kimi-K2-Instruct-0905">https://huggingface.co/moonshot-ai/Kimi-K2-Instruct-0905</a>	1040B MoE
Gemma 3 VL	<a href="https://huggingface.co/google/gemma-3-12b-it">https://huggingface.co/google/gemma-3-12b-it</a>	12B
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-32B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-32B-Instruct</a>	32B
DeepSeek V3	<a href="https://huggingface.co/deepseek-ai/DeepSeek-V3-0324">https://huggingface.co/deepseek-ai/DeepSeek-V3-0324</a>	671B MoE
SEA-LION v3 (Llama 3.1)	<a href="https://huggingface.co/aisingapore/Llama-SEA-LION-v3-70B-IT">https://huggingface.co/aisingapore/Llama-SEA-LION-v3-70B-IT</a>	70B
Tulu 3	<a href="https://huggingface.co/allenai/Llama-3.1-Tulu-3-70B">https://huggingface.co/allenai/Llama-3.1-Tulu-3-70B</a>	70B
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-14B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-14B-Instruct</a>	14B
SEA-LION v4 (Qwen 3 VL)	<a href="https://huggingface.co/aisingapore/Qwen-SEA-LION-v4-8B-VL">https://huggingface.co/aisingapore/Qwen-SEA-LION-v4-8B-VL</a>	8B
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct</a>	8B
Qwen 2.5	<a href="https://huggingface.co/Qwen/Qwen2.5-72B-Instruct">https://huggingface.co/Qwen/Qwen2.5-72B-Instruct</a>	72B
Qwen 2.5	<a href="https://huggingface.co/Qwen/Qwen2.5-32B-Instruct">https://huggingface.co/Qwen/Qwen2.5-32B-Instruct</a>	32B
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-8B-Instruct</a>	8B
Mistral Large 2411	<a href="https://huggingface.co/mistralai/Mistral-Large-Instruct-2411">https://huggingface.co/mistralai/Mistral-Large-Instruct-2411</a>	123B
Qwen 3 A3B	<a href="https://huggingface.co/Qwen/Qwen3-30B-A3B">https://huggingface.co/Qwen/Qwen3-30B-A3B</a>	30B MoE
Gemma 2	<a href="https://huggingface.co/google/gemma-2-27b-it">https://huggingface.co/google/gemma-2-27b-it</a>	27B
SEA-LION v4 (Qwen 3 VL)	<a href="https://huggingface.co/aisingapore/Qwen-VL-SEA-LION-v4">https://huggingface.co/aisingapore/Qwen-VL-SEA-LION-v4</a>	4B
Llama 3.3	<a href="https://huggingface.co/meta-llama/Llama-3.3">https://huggingface.co/meta-llama/Llama-3.3</a>	70B
Llama 3.1	<a href="https://huggingface.co/meta-llama/Llama-3.1">https://huggingface.co/meta-llama/Llama-3.1</a>	70B
SEA-LION v3 (Llama 3.1)	<a href="https://huggingface.co/aisingapore/Llama-SEA-LION-v3-8B-IT">https://huggingface.co/aisingapore/Llama-SEA-LION-v3-8B-IT</a>	8B
ERNIE 4.5	<a href="https://huggingface.co/ERNIE/ERNIE-4.5">https://huggingface.co/ERNIE/ERNIE-4.5</a>	21B MoE
Command A 03-2025	<a href="https://huggingface.co/CohereLabs/c4ai-command-a-03-2025">https://huggingface.co/CohereLabs/c4ai-command-a-03-2025</a>	111B
SEA-LION v3 (Gemma 2)	<a href="https://huggingface.co/aisingapore/Gemma-SEA-LION-v3-9B-IT">https://huggingface.co/aisingapore/Gemma-SEA-LION-v3-9B-IT</a>	9B
Qwen 2.5	<a href="https://huggingface.co/Qwen/Qwen2.5-14B-Instruct">https://huggingface.co/Qwen/Qwen2.5-14B-Instruct</a>	14B
Llama 3	<a href="https://huggingface.co/meta-llama/Meta-Llama-3-70B">https://huggingface.co/meta-llama/Meta-Llama-3-70B</a>	70B
Apertus	<a href="https://huggingface.co/swiss-ai/Apertus-70B-2509">https://huggingface.co/swiss-ai/Apertus-70B-2509</a>	70B
Sailor2	<a href="https://huggingface.co/sail/Sailor2-8B">https://huggingface.co/sail/Sailor2-8B</a>	8B
Tulu 3	<a href="https://huggingface.co/allenai/Llama-3.1-Tulu-3-8B-SFT">https://huggingface.co/allenai/Llama-3.1-Tulu-3-8B-SFT</a>	8B
MERaLION 2	<a href="https://huggingface.co/MERaLION/MERaLION-2-10B">https://huggingface.co/MERaLION/MERaLION-2-10B</a>	10B
Babel	<a href="https://huggingface.co/Tower-Babel/Babel-83B">https://huggingface.co/Tower-Babel/Babel-83B</a>	83B
Gemma 2	<a href="https://huggingface.co/google/gemma-2-9b-it">https://huggingface.co/google/gemma-2-9b-it</a>	9B
Apertus	<a href="https://huggingface.co/swiss-ai/Apertus-8B-2509">https://huggingface.co/swiss-ai/Apertus-8B-2509</a>	8B
Sailor2	<a href="https://huggingface.co/sail/Sailor2-20B">https://huggingface.co/sail/Sailor2-20B</a>	20B
Llama 3.1	<a href="https://huggingface.co/meta-llama/Llama-3.1-8B">https://huggingface.co/meta-llama/Llama-3.1-8B</a>	8B
Qwen 2.5	<a href="https://huggingface.co/Qwen/Qwen2.5-7B-Instruct">https://huggingface.co/Qwen/Qwen2.5-7B-Instruct</a>	7B
Babel	<a href="https://huggingface.co/Tower-Babel/Babel-9B">https://huggingface.co/Tower-Babel/Babel-9B</a>	9B
SeaLLMs V3	<a href="https://huggingface.co/SeaLLMs/SeaLLMs-v3-7B-Chat">https://huggingface.co/SeaLLMs/SeaLLMs-v3-7B-Chat</a>	7B
Command R+ 08-2024	<a href="https://huggingface.co/CohereLabs/c4ai-command-r-plus-08-2024">https://huggingface.co/CohereLabs/c4ai-command-r-plus-08-2024</a>	104B
phi-4	<a href="https://huggingface.co/microsoft/phi-4">https://huggingface.co/microsoft/phi-4</a>	14B
Aya Expans	<a href="https://huggingface.co/CohereLabs/aya-expans-32b">https://huggingface.co/CohereLabs/aya-expans-32b</a>	32B
Command R 08-2024	<a href="https://huggingface.co/CohereLabs/c4ai-command-r-08-2024">https://huggingface.co/CohereLabs/c4ai-command-r-08-2024</a>	32B
Olmo 2 0325	<a href="https://huggingface.co/allenai/OLMo-2-0325-32B-Instruct">https://huggingface.co/allenai/OLMo-2-0325-32B-Instruct</a>	32B
Minstral 2410	<a href="https://huggingface.co/mistralai/Minstral-8B-Instruct-2410">https://huggingface.co/mistralai/Minstral-8B-Instruct-2410</a>	8B
Olmo 3	<a href="https://huggingface.co/allenai/OLmo-3-7B-Instruct">https://huggingface.co/allenai/OLmo-3-7B-Instruct</a>	7B
Llama 3	<a href="https://huggingface.co/meta-llama/Meta-Llama-3-8B">https://huggingface.co/meta-llama/Meta-Llama-3-8B</a>	8B
Command R7B 12-2024	<a href="https://huggingface.co/CohereLabs/c4ai-command-r7b-12-2024">https://huggingface.co/CohereLabs/c4ai-command-r7b-12-2024</a>	7B
Aya Expans	<a href="https://huggingface.co/CohereLabs/aya-expans-8b">https://huggingface.co/CohereLabs/aya-expans-8b</a>	8B
Mistral Small 3.1 2503	<a href="https://huggingface.co/mistralai/Mistral-Small-3.1-24B-Instruct-2503">https://huggingface.co/mistralai/Mistral-Small-3.1-24B-Instruct-2503</a>	24B
Olmo 2 1124	<a href="https://huggingface.co/allenai/OLMo-2-1124-7B-Instruct">https://huggingface.co/allenai/OLMo-2-1124-7B-Instruct</a>	7B
Olmo 2 1124	<a href="https://huggingface.co/allenai/OLMo-2-1124-13B-Instruct">https://huggingface.co/allenai/OLMo-2-1124-13B-Instruct</a>	13B
SEA-LION v4 (Gemma 3 VL)	<a href="https://huggingface.co/aisingapore/Gemma-SEA-LION-v4-4B-VL">https://huggingface.co/aisingapore/Gemma-SEA-LION-v4-4B-VL</a>	4B
Gemma 3 VL	<a href="https://huggingface.co/google/gemma-3-4b-it">https://huggingface.co/google/gemma-3-4b-it</a>	4B
Qwen 3 VL	<a href="https://huggingface.co/Qwen/Qwen3-VL-4B-Instruct">https://huggingface.co/Qwen/Qwen3-VL-4B-Instruct</a>	4B
SEA-LION v4 (Apertus)	<a href="https://huggingface.co/aisingapore/Apertus-SEA-LION-v4-8B-IT">https://huggingface.co/aisingapore/Apertus-SEA-LION-v4-8B-IT</a>	8B
<i>Reasoning Models</i>		
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-235B-A22B-Thinking-2507">https://huggingface.co/Qwen/Qwen3-235B-A22B-Thinking-2507</a>	235B MoE
DeepSeek V3.1 Thinking	<a href="https://huggingface.co/deepseek-ai/DeepSeek-V3.1">https://huggingface.co/deepseek-ai/DeepSeek-V3.1</a>	671B MoE
Qwen 3 Next (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-Next-80B-A3B-Thinking">https://huggingface.co/Qwen/Qwen3-Next-80B-A3B-Thinking</a>	80B MoE
Deepseek R1 0528	<a href="https://huggingface.co/deepseek-ai/DeepSeek-R1-0528">https://huggingface.co/deepseek-ai/DeepSeek-R1-0528</a>	671B MoE
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-30B-A3B-Thinking-2507">https://huggingface.co/Qwen/Qwen3-30B-A3B-Thinking-2507</a>	30B MoE
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-32B">https://huggingface.co/Qwen/Qwen3-32B</a>	32B
GPT OSS	<a href="https://platform.openai.com/docs/models">https://platform.openai.com/docs/models</a>	120B MoE mxfp4
SEA-LION v3.5 R (Llama)	<a href="https://huggingface.co/aisingapore/Llama-SEA-LION-v3.5-70B-R">https://huggingface.co/aisingapore/Llama-SEA-LION-v3.5-70B-R</a>	70B
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-14B">https://huggingface.co/Qwen/Qwen3-14B</a>	14B
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3-8B">https://huggingface.co/Qwen/Qwen3-8B</a>	8B
GPT OSS	<a href="https://huggingface.co/openai/gpt-oss-20b">https://huggingface.co/openai/gpt-oss-20b</a>	20B MoE mxfp4
Qwen 3 (Thinking)	<a href="https://huggingface.co/Qwen/Qwen3">https://huggingface.co/Qwen/Qwen3</a>	4B
QwQ	<a href="https://huggingface.co/Qwen/QwQ-32B-Preview">https://huggingface.co/Qwen/QwQ-32B-Preview</a>	32B
Reka Flash 3.1	<a href="https://huggingface.co/RekaAI/reka-flash-3.1">https://huggingface.co/RekaAI/reka-flash-3.1</a>	21B
SEA-LION v3.5 R (Llama)	<a href="https://huggingface.co/aisingapore/Llama-SEA-LION-v3.5-8B-R">https://huggingface.co/aisingapore/Llama-SEA-LION-v3.5-8B-R</a>	8B
Olmo 3 Think	<a href="https://huggingface.co/allenai/OLmo-3-32B-Think">https://huggingface.co/allenai/OLmo-3-32B-Think</a>	32B
Olmo 3 Think	<a href="https://huggingface.co/allenai/OLmo-3-7B-Think">https://huggingface.co/allenai/OLmo-3-7B-Think</a>	7B

Table 9: Links and sizes of models evaluated in our current experiments.