

Multilingual LLMs Inherently Reward In-Language Time–Sensitive Semantic Alignment for Low-Resource Languages

Ashutosh Bajpai^{1,2}, Tanmoy Chakraborty¹

¹ Indian Institute of Technology Delhi, India

² Wipro Research, India

eez228482@ee.iitd.ac.in, tanchak@ee.iitd.ac.in

Abstract

The unwavering disparity in labeled resources between resource-rich languages and those considered low-resource remains a significant impediment for Large Language Models (LLMs). Recent strides in cross-lingual in-context learning (X-ICL), mainly through semantically aligned examples retrieved from multilingual pre-trained transformers, have shown promise in mitigating this issue. However, our investigation reveals that LLMs intrinsically reward in-language semantically aligned cross-lingual instances over direct cross-lingual semantic alignments, with a pronounced disparity in handling time–sensitive queries in the X-ICL setup. Such queries demand sound temporal reasoning ability from LLMs, yet the advancements have predominantly focused on English. This study aims to bridge this gap by improving temporal reasoning capabilities in low-resource languages. To this end, we introduce `mTEMPREASON`, a temporal reasoning dataset aimed at the varied degrees of low-resource languages and propose **Cross-Lingual Time-Sensitive Semantic Alignment (CLiTSSA)**, a novel method to improve temporal reasoning in these contexts. To facilitate this, we construct an extension of `mTEMPREASON` comprising pairs of parallel cross-language temporal queries along with their anticipated in-language semantic similarity scores. Our empirical evidence underscores the superior performance of `CLiTSSA` compared to established baselines across three languages – Romanian, German, and French, encompassing three temporal tasks and including a diverse set of four contemporaneous LLMs. This marks a significant step forward in addressing resource disparity in the context of temporal reasoning across languages.

Introduction

In the evolving landscape of Large Language Models (LLMs), temporal reasoning requires models to comprehend and interpret the significant subtleties inherent in time–time, time–event and event–event correlations (Chen, Wang, and Wang 2021; Dhingra et al. 2022). Temporality is a crucial dimension of information that evolves through creation, maintenance, and obsolescence. Enhancing LLMs with this faculty augments their analytical capabilities, paving the way for addressing intricate challenges prevalent in domains sensitive to temporal dynamics, such as finance, healthcare, le-

gal studies, and archaeology. Furthermore, addressing low-resource languages in LLMs is crucial for computational linguistics, given their paucity of data and digital infrastructure (Cahyawijaya et al. 2023; Asai et al. 2023; Adilazuarda et al. 2024). Enhancing LLMs for these languages improves not just genuine linguistic inclusivity but also their application and acceptance across diverse cultural landscapes. The discourse on enhancing temporal reasoning in LLMs has, until now, been predominantly focused on English. Our work seeks to alleviate this disparity by propelling temporal reasoning in low-resource languages.

Cross-Lingual In-Context Prompting. Recent advancements in in-context learning (ICL), prompted by the advent of LLMs, have shown promising results (Zhao et al. 2021; Lin et al. 2022b; Liu et al. 2022; Zhang et al. 2022). The stark contrast in annotated data availability among languages accentuates the usage of high-resource linguistic contexts for addressing tasks in low-resource languages. The ICL approach was adapted by Winata et al. (2021) for cross-lingual (X-ICL) applications by randomly selecting examples from a resource-rich language to support queries in a language with limited resources.

Figure 1 shows three levels of temporal queries - L1 (Time–Time), L2 (Time–Event), and L3 (Event–Event) and the expected model responses: time information for L1 queries, an event for L2 given time, and for L3, an event in response to another event, without explicit temporal details in either input or output. An illustrative L1 query in English (“*What is the time 6 years and 4 months after Nov, 1185*”) with an associated answer (“*Mar, 1192*”) serves as additional context for a corresponding L1 query in Romanian (“*Care este timpul cu 8 ani și 3 luni înainte de august 1240*”). Subsequent research suggested that cross-lingual examples that are semantically aligned can significantly enhance performance compared to the arbitrarily selected examples (Tanwar et al. 2023). Further development also indicates that semantic similarity alone doesn’t ensure the optimal performance, stressing the necessity for a learning-based retrieval model. (Lin, Martins, and Schütze 2024).

Challenges in Cross-Lingual Semantic Alignment. Cross-lingual in-context approaches rely on the contextual semantic profoundness embedded within multilingual pre-trained encoder-only transformers, reflected through

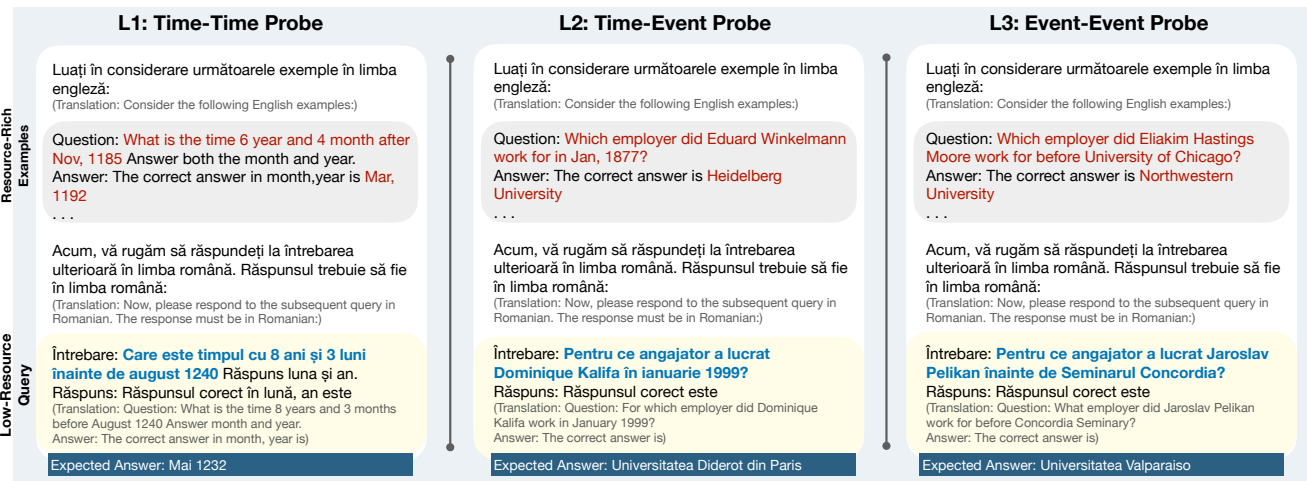


Figure 1: A working example of the low-resource cross-lingual prompting across three temporal tasks: L1, L2, and L3 in Romanian. The translations included in small brackets are not integral to the prompt; their purpose is solely to enhance readability.

their embedding space, for retrieving semantically akin examples. Nonetheless, the disparity in linguistic distribution within the pre-training dataset, favoring resource-rich languages over those with fewer resources, significantly hinders efficient cross-lingual semantic alignment within their embedding space, especially concerning time-sensitive queries. The ensuing analysis is conducted to validate the aforementioned hypothesis in a temporal context.

The objective is to assess the efficacy of multilingual sentence-BERT (Reimers and Gurevych 2019) in retrieving semantically akin cross-lingual in-context examples for L1 task, considering Romanian, German, and French as low and English as a high-resource language. The investigation juxtaposes two distinct approaches: firstly, *cross-lingual-similarity*, wherein the top-3 English instances are directly retrieved for a low-resource query by ranking the similarity scores between a low-resource query and English instances. Secondly, *in-language-similarity*, where this is achieved by initially translating the English example dataset into low-resource languages. Next, the top-3 instances are sourced using similarity scores between a low-resource query and the translated English examples dataset. Subsequently, the translated examples are replaced with their corresponding English instances, thus retrieving the English examples with the highest in-language similarity. Performances are quantified using F1 scores and exact match (EM) scores. The results presented in Table 1 indicate that multilingual encoder-only transformers exhibit profound semantic similarity for temporal queries in an in-language similarity framework, outperforming their counterparts in a cross-lingual similarity context when identifying semantically akin examples for the cross-lingual temporal reasoning task across languages. Consequently, this underscores the need to evaluate and enhance the cross-lingual, time-sensitive semantic context of retrieval models. Nonetheless, the availability of data for such alignment in low-density languages is limited.

Our Proposed Method. We start with the development of first-of-its-kind a comprehensive benchmark dataset,

Task	Settings	F1	EM
Romanian	Cross-Lingual-Similarity	33.65	10.15
	In-Language-Similarity	43.48	18.65
German	Cross-Lingual-Similarity	56.63	35.45
	In-Language-Similarity	64.77	46.35
French	Cross-Lingual-Similarity	46.62	22.05
	In-Language-Similarity	60.50	37.17

Table 1: Analyzing the impact of in-language versus cross-language similarities in retrieving semantically similar akin examples in a three-shot cross-lingual setup for the L1 task across languages using LLaMA3-8B

mTEMPREASON, to evaluate temporal reasoning for limited resource languages – Romanian, German, and French across a diverse set of LLMs. Further, we endeavor to devise an efficacious novel cross-lingual retriever, CLITSSA (Cross-Lingual Time-Sensitive Semantic Alignment), for handling time-sensitive queries from low-resource languages, addressing the aforementioned challenges in a cross-lingual context. To achieve this, drawing inspiration from Yamada and Ri (2024), we elect to apply the transfer of profound semantic space knowledge within a language to a cross-lingual semantic space for queries influenced by temporality. Consequently, we adopt a supervised fine-tuning approach that necessitates an additional dataset to facilitate this transition. To this end, we curate an extension of the mTEMPREASON dataset comprising parallel sentences for Romanian-English, German-English, and French-English pairs, accompanied by their anticipated similarity scores in the semantic space of the low-resource language. By employing this curated dataset, the transition of the semantic context from a monolingual to a cross-lingual embedding sphere is attained.

Remarkably, for temporal queries, CLITSSA outperforms the arbitrary cross-lingual in-context benchmark, demonstrating relative mean F1 score enhancements of 11.41%, 30.77%, and 62.92% for Romanian, German, and French, respectively. Additionally, it evidences a significant

	Train	Dev	Test
Time Range	1014-2022	634-2023	998-2023
L1	400,000	4,000	4,000
L2	16,017	5,521	5,397
L3	13,014	4,437	4,426

Table 2: Dataset statistics for mTEMPREASON.

improvement in relative mean F1 score of 6.38%, 5.98%, and 20.93% compared to the contemporary cross-lingual in-context baselines for Romanian, German, and French.

Our contributions are summarized below¹–

- We develop a dataset centered around temporal reasoning, mTEMPREASON, elusively designed for varying degrees of limited resource languages.
- Our findings reveal that multilingual transformers exhibit superior in-language semantic similarity over cross-lingual similarity context for temporal queries, especially explicit ones, in the X-ICL setup.
- We introduce CLITSSA to enhance temporal reasoning capabilities within LLMs for low-resource languages. Consequently, we develop an extension of mTEMPREASON comprised of paired cross-lingual time-sensitive queries with corresponding similarity scores. Our empirical analysis demonstrates that CLITSSA significantly outperforms the contemporary baselines.

Benchmark For Low-Resource Temporal Reasoning

In our study of temporal reasoning, the TEMPREASON (Tan, Ng, and Bing 2023) stands out as a recent, comprehensive resource, providing multifaceted temporality across an extended time frame. Therefore, we select it to develop the first multilingual, low-resource dataset for temporal reasoning. To this end, we employ the T5 model (Raffel et al. 2023) to automatically translate the dataset from English language to Romanian, German and French languages.

The mTEMPREASON Dataset

TEMPREASON encompasses tasks categorized into three levels of temporal complexity, namely time-time, time-event, and event-event relationships, corresponding to Levels L1, L2, and L3, respectively. The dataset’s statistical information, along with the partition into training (Train), development (Dev), and test set (Test), is detailed in Table 2. Employing a concise prompt prefix, “Translate the following sentences from English to the {Target Language},” we employ the T5 model to translate this dataset into a selection of languages with varying degrees of limited resources, specifically French, German, and Romanian. We opt for these three languages in this work as they provide varied levels of limited resources compared to English.

¹Source code and dataset are available at <https://github.com/ab-ittd/clitssa>.

Metric	Setting	Ro.	Ge.	Fr.	Avg.
TSR	Auto	99.00	97.57	97.79	98.11 ± 1.8
BTA	BLEU-3	51.12	56.47	43.63	50.41

Table 3: Quality assessment of mTEMPREASON’s translations: employing automated verification for *Translation Success Rate (TSR in %)*, and applying BLEU-3 and manual review standards for *Back-Translation Accuracy (BTA)* evaluation across languages— Romanian (Ro.), German (Ge.) and French(Fr.), averaging over temporal tasks.

Romanian, German, and French have 98.3%, 89.5%, and 78% fewer speakers compared to English², respectively.

Data Quality

The mTEMPREASON dataset was constructed by a linguist who specialized in NLP³. We employed back-translation-based (Miyabe and Yoshino 2015) evaluation to ensure the superior quality of the proposed dataset. A random selection of 100 query examples was made from mTEMPREASON for each of the translated languages — Romanian, German, and French, across temporal tasks within the test dataset. These queries underwent back-translation⁴ into the source language (English) and were subsequently compared to their original counterparts in the resource-rich language (English) to assess fidelity and coherence. The analysis employed the BLEU-3 metric for quantitative evaluation. In addition, successful translation from the source to the target languages was noted, regardless of translation quality. The translation success rate was documented using an automated approach by employing a language detection library⁵ across the entire test dataset. In Table 3, the mean automated translation success rate was recorded at 98.11 ± 1.83%. Concurrently, a mean BLEU-3 score of 50.41 was observed for assessing back-translation-based accuracy.

Problem Setting

The Time-Sensitive Question Answering (TSQA) task requires that the LLMs generate an accurate answer in response to a temporal query. This answer may be a temporal delineation or an event, depending on the structure of the query, which can include time-time, time-event, and event-event scenarios. Our experiments are conducted in a closed-book environment, requiring LLMs to deliver precise facts without reliance on external contexts. As illustrated in Figure 1, to demonstrate the prompt in a cross-lingual in-context learning framework, a few-shot example from a resource-rich language, accompanied by the query in a low-resource language, is provided.

In this study, we consider the following baselines–

- **Cross-Lingual In-Context Learning (X-ICL) (Winata et al. 2021)**. The model is primed with limited examples

²<https://www.ethnologue.com/insights/ethnologue200>

³A male expert, within the 30-40 age bracket.

⁴<https://translate.google.com>

⁵<https://pypi.org/project/langdetect>

from resource-rich language serving as demonstrations, along with a query in low-resource language.

- **X-InSTA (Semantic Aligner) (Tanwar et al. 2023).** X-InSTA has advanced the X-ICL method by introducing a retrieval of semantically akin examples for queries across languages, leveraging label space alignment.

Method

Primer

Let us consider a resource-rich source dataset D_r , containing pairs of queries and answers (q_i^r, a_i^r) for each $i \in m$, with m representing the total number of samples within D_r . Additionally, let D_l be a low-resource language dataset, which similarly comprises query and answer pairs (q_j^l, a_j^l) for each $j \in n$, where n signifies the total sample count in D_l . Within a conventional ICL framework, K arbitrary instances of question-answer pairs are selected from D_r , designated as context C for a low-resource query q_x^l from D_l , with $x \in n$. The goal is to optimize the expected value of a_x^l , given context C and the query input as illustrated in Equation 1, where A^l represents the vocabulary space corresponding to query $q_x^l \in D_l$.

$$a_x^l = \arg \max_{a^l \in A^l} p(a^l | C, q_x^l) \quad (1)$$

In the case of a semantic aligner, C is constructed to maximize the semantic alignment between query q_x^l and context C . Hence, we introduce $e_{q_i^r}$ as a dense embedding representation produced by a multilingual pre-trained transformer for query q_i^r . Correspondingly, for each $i \in m$, $e_{q_i^r}$ represents the embeddings for queries within a resource-rich dataset D_r . Furthermore, $f(s)$ and $f(d)$ denote similarity and distance functions, respectively, such that for cosine similarity, $f(s) = 1 - f(d)$; a lesser distance implies greater similarity. The overarching goal is to identify a set of K examples, denoted as S_K , where the semantic similarity surpasses that of other dataset examples, as shown in Equation 2.

$$S_K = \{(q_k^r, a_k^r) \forall k \in \{1, \dots, K\}, \\ \text{if } f(s)_{q_k^r, q_x^l} \geq f(s)_{q_z^r, q_x^l} \forall z \in m \text{ and } z \notin K\} \quad (2)$$

The variable $f(s)_{q_k^r, q_x^l}$ denotes the semantic similarity between a low-resource query q_x^l and an example query q_k^r from a resource-rich language. To calculate $f(s)_{q_k^r, q_x^l}$, the procedure starts with the extraction of dense embeddings, $e_{q_k^r}$ and $e_{q_x^l}$ for the input queries q_k^r and q_x^l , respectively. This extraction is performed utilizing a multilingual pre-trained transformer model. Subsequently, the function $f(d)$ is applied to these embeddings to measure the distance, which ultimately yields the value of $f(s)$.

Cross-Lingual Time-Sensitive Semantic Alignment (CLiTSSA)

Objective. We introduce CLiTSSA to augment the semantic similarity context of time-sensitive queries within the cross-lingual embedding space. We elect to employ the transfer of the comprehensive time-sensitive semantic knowledge from an in-language embedding space to a

cross-lingual embedding space. For this purpose, we have embraced a supervised fine-tuning approach, making use of a training dataset comprised of sentence pairs along with their associated labeled scores, quantifying the expected time-sensitive semantic similarity among pairs of queries. The objective is to achieve an effective cross-lingual time-sensitive contextual alignment of temporal queries for LLMs, thus boosting in-context performance.

Training Dataset. A training dataset D_t is constructed comprising pairs of sentences alongside their associated similarity scores. Specifically, D_t consists of $(q_u^l, q_v^r | f(s)_{u,v})$, where q_u^l denotes a low-resource query derived from D_l with $u \in n$, q_v^r indicates a query from a resource-rich dataset D_r with $v \in m$, and $f(s)_{u,v}$ represents the similarity score between these queries within a low-resource monolingual embedding space.

Here, we present a systematic approach to construct D_t . Initially, we delineate D'_r , a transformed resource-rich dataset in a low-resource language using translation. Subsequently, we determine the temporal semantic alignment scores $f(s)$ among the queries in D_l and D'_r . The utilization of all example pairs in the fine-tuning procedure incurs quadratic complexity in terms of $|D_t|$, rendering it resource-intensive. Drawing inspiration from Rubin, Herzig, and Berant (2022), this issue is addressed by selecting the top- h analogous examples from D'_r for each query in D_l . To counteract training data bias due to high similarity scores, we randomly select w examples from the remaining dataset to capture the whole similarity distribution. Consequently, for every query $q_u^l \in D_l$, the resultant set S_u comprises $(h+w)$ sentence pairs, each accompanied by their temporal semantic similarity scores as postulated in Equation 3.

$$S_u = \{(q_u^l, q_1^{r'} | f(s)_{u,1}), \dots, (q_u^l, q_h^{r'} | f(s)_{u,h}), \\ (q_u^l, q_{h+1}^{r'} | f(s)_{u,h+1}), \dots, (q_u^l, q_{h+w}^{r'} | f(s)_{u,h+w})\} \quad (3)$$

Likewise, paired sentences and their associated similarity scores are generated for all queries within D_l . In the concluding phase, the transformed, resource-rich dataset D'_r is substituted back with the original dataset D_r . Specifically, the query $q_v^{r'}$, which represents a transformation of the resource-rich query q_v^r into a low-resource language, is replaced with the original query q_v^r within the paired sentences dataset. Multilingual Sentence-BERT (Reimers and Gurevych 2019), a pre-trained transformer model, is employed to derive the semantic alignment scores.

Fine-tuning the Retriever. CoSENT⁶ (Cosine Sentence) loss is employed for the fine-tuning of sentence pairs along with similarity scores as labels, utilizing multilingual Sentence-BERT as base retriever. CoSENT loss generates a more robust training signal for optimizing the cosine value than the traditional cosine similarity loss function. This loss function is shown in Equation 4.

$$\mathcal{L} = \log \sum (1 + \exp(f(s)_{(q_a^l, q_b^r)} - f(s)_{(q_y^l, q_z^r)} + \exp \dots)) \quad (4)$$

⁶<https://kexue.fm/archives/8847>

Task	Method	F1.				EM			
		French	German	Romanian	Avg.	French	German	Romanian	Avg.
L1	X-ICL	33.60	45.33	34.17	37.70	14.85	22.79	08.80	15.48
	X-InSTA [†]	46.62	56.63	33.65	45.63	22.05	35.45	10.15	22.55
	CLiTSSA	57.15	59.77	37.16	51.36	32.57	39.3	13.45	28.44
	CLiTSSA*	55.04	63.50	37.78	52.11	31.62	45.15	13.70	30.16
L2	X-ICL	11.00	08.96	10.99	10.32	03.57	03.73	03.97	03.76
	X-InSTA [†]	11.92	12.45	11.40	11.92	04.55	05.18	03.83	04.52
	CLiTSSA	15.23	14.02	11.42	13.56	05.81	05.24	03.87	04.97
	CLiTSSA*	15.04	13.53	11.71	13.43	05.74	05.20	04.03	04.99
L3	X-ICL	17.07	16.18	18.84	17.36	07.28	05.77	08.82	07.29
	X-InSTA [†]	17.74	18.35	22.18	19.42	10.55	09.08	12.85	10.83
	CLiTSSA	19.87	18.52	22.94	20.44	11.22	08.85	13.33	11.13
	CLiTSSA*	19.92	18.69	22.53	20.38	11.45	08.72	13.23	11.13
	$\overline{\Delta}_{\text{CLiTSSA} \rightarrow \uparrow}$	5.32 \uparrow	1.62 \uparrow	1.43 \uparrow	2.79 \uparrow	4.15 \uparrow	1.23 \uparrow	1.27 \uparrow	2.22 \uparrow
	$\overline{\Delta}_{\text{CLiTSSA}^{max} \rightarrow \uparrow}$	5.34 \uparrow	2.92 \uparrow	1.73 \uparrow	3.04 \uparrow	4.22 \uparrow	3.17 \uparrow	1.41 \uparrow	2.79 \uparrow

Table 4: Comparison of F1 and EM (Exact Match) scores across different prompting strategies for temporal tasks and languages in a three-shot setup employing LLaMA3-8B. The strategies include *X-ICL* and *X-InSTA*, representing random and semantically aligned cross-lingual baselines, respectively, while CLiTSSA* denotes an integrated retriever trained across languages and tasks, and CLiTSSA indicates a language and task-specific retriever. $\overline{\Delta}$ represents mean improvement for languages across temporal tasks and CLiTSSA^{max} representing $\max(\text{CLiTSSA}, \text{CLiTSSA}^*)$. We report mean values over three runs by varying the parameter $top-p \in \{1.0, 0.8, 0.6\}$ and apply one tail Mann-Whitney U test for p-values. We observe a p-value of 0.05 while comparing the mean F1 score of CLiTSSA with X-InSTA across languages and tasks.

In this context, where (q_a^l, q_b^r) and (q_y^l, q_z^r) represent instances from D_t within a batch, under the condition that the anticipated similarity between (a, b) exceeds that of (y, z) , the summation extends across all feasible input pairs within the batch satisfying this criterion. This approach amalgamates both cross-entropy and contrastive loss advantages.

Experimental Results And Analysis

Experimental Setup

Primarily, we employ LLaMA3-8B (AI@Meta 2024) for all experimental works. A three-shot ICL approach is used throughout the experimental setting, demonstrating superior outcomes compared to both one-shot and two-shot configurations. The value of h and w is set empirically at 30 and 10, respectively. To fine-tune the CLiTSSA retriever model, the ‘distiluse-base-multilingual-cased-v1’ serves as the foundational model. This method is systematically applied to each low-resource language across temporal tasks – L1, L2 and L3, to ensure optimum performance. Additionally, an integrated CLiTSSA retriever is fine-tuned across languages and temporal tasks. The Train and Dev datasets from mTEMPREASON are used to construct the parallel corpus to fine-tune the CLiTSSA retriever, with a separate held-out test set employed to benchmark all outcomes. We use word level F1 scores and exact match (EM) standards to quantify the LLM’s responses. Please refer to the technical appendix for ablations on few-shots, parameters h & w , along with hyperparameters in detail.

CLiTSSA Advancements Over Precedence

The comprehensive comparison of CLiTSSA with baselines highlights the effectiveness of incorporating cross-lingual

Model	Method	F1.		
		L1	L2	L3
LLaMA3-8B	X-InSTA [†]	46.62	11.92	17.74
	CLiTSSA	57.15	15.23	19.87
Mistral-v1	X-InSTA [†]	38.26	11.73	18.72
	CLiTSSA	46.45	14.83	18.64
Vicuna-7b-v1.5	X-InSTA [†]	27.93	9.54	12.68
	CLiTSSA	36.67	12.04	12.89
Bloomz-7b1	X-InSTA [†]	29.45	3.79	13.49
	CLiTSSA	40.74	4.20	13.57
	$\overline{\Delta}_{\text{CLiTSSA} \rightarrow \uparrow}$	9.68 \uparrow	2.33 \uparrow	0.58 \uparrow

Table 5: The performance of CLiTSSA across LLMs for temporal tasks using the French test set ($\overline{\Delta}$: the mean improvement in F1 score across LLMs for a temporal task).

time-sensitive semantic alignment compared to a conventional semantic aligner (X-InSTA), as evidenced in Table 4 across a variety of low-resource languages and temporal tasks. The mean values of metrics were compared across three iterations by varying the model’s parameter $top-p$ (1.0, 0.8, 0.6). The parameter indicates the cumulative probability threshold for token selection. Notably, CLiTSSA achieves a mean increase of 5.32, 1.62, and 1.43 points in the F1-score for French, German, and Romanian, respectively, with a p-value of 0.05. Specifically, the most significant improvements in F1-score—10.53, 3.31, and 2.13 points for tasks L1, L2, and L3, respectively are observed in the French setting. A similar enhancement is discernible concerning the EM metric. Moreover, the overall analysis does not yield a definitive conclusion for an integrated CLiTSSA retriever




Figure 2: Comparison of F1 scores using box plot: a dual perspective on temporal tasks with language models and languages, pivoting on French and LLaMA3-8B, respectively.

over its counterpart except a notable transcend of 3.7 points in the F1-score for task L1 within the German setting.

Robustness Across LLMs

The assessment of CLiTSSA across a variety of distinct, contemporary LLMs namely, an English-dominant instruction-tuned model, Vicuna (Zheng et al. 2023), a model fluent in French and German, Mistral (Jiang et al. 2023), and a cross-lingual specialized LLM, Bloomz (Muenighoff et al. 2023), demonstrates the robustness of the method. This evaluation highlights the versatility and effectiveness of CLiTSSA in engaging with and analyzing linguistic data across different languages and model architectures. The findings, as detailed in Table 5, reveal that CLiTSSA surpasses the baseline by margins of 9.68, 2.33, and 0.58 points in mean F1-score for L1, L2, and L3, respectively, across LLMs.

To further substantiate the statistical significance of findings, a comparison is drawn through a box plot analysis of the F1-scores under two scenarios: firstly, by plotting F1-scores across LLMs and temporal tasks with a focus on the French language and secondly, through a box plot that contrasts F1-scores across various languages and temporal tasks centered around a specific LLM, namely LLaMA3-8B. Figure 2 shows that CLiTSSA notably extends the upper quartile by 10.01 and 5.26 points in terms of F1-score, with a mean increase of 4.20 and 2.79 points in the F1 scores, across these scenarios, respectively. Additionally, embedding space evolution under CLiTSSA, presented in technical appendix, further elucidates the noted enhancement.

Cross-Task CLiTSSA Performance

Here, we evaluate the CLiTSSA retriever’s generalization across temporal tasks to see if semantic alignment, sensitive to temporal variation achieved in one task, facilitates the resolution of another temporal task without necessitating a re-fine-tuning. The CLiTSSA model, once fine-tuned on a specific temporal task, is assessed on two other temporal tasks, i.e., the retriever optimized on the L1 task is employed to retrieve time-sensitive semantic examples for the L2 and L3 tasks. Figure 3 presents the outcomes of this investigation. Note that tasks L1, L2, and L3 are sequentially arranged in order of temporal complexity, with L3 being the most intricate. The findings reveal that the temporal alignment ac-




Figure 3: Cross-Task CLiTSSA performance across tasks with F1 scores on the French test set against the X-InSTA baseline. CLiTSSA-L* represents a retriever fine-tuned using L* training dataset where $* \in \{1, 2, 3\}$.

quired through the lower-level temporal task (L1) can significantly enhance the relative F1-score of the more complex tasks L2 and L3 by 13.5% and 14.0%, respectively. However, the reverse scenario is inapplicable. Moreover, more complex tasks, L2 and L3, can exchange learning, thereby improving their F1-score relatively by 27.1% and 13.6%, respectively. French is employed as the low-resource language in this study. The results corroborate that fine-tuning the CLiTSSA with a low-level temporal task (L1) could serve as a superior alternative to any semantic-based example retriever across temporal tasks.

Cross-Linguality vs. Monolinguality

The complexity of the prompt increases with the incorporation of multiple languages, which detrimentally impacts the performance of ICL in cross-lingual contexts when contrasted with monolingual scenarios. This experiment delineates the CLiTSSA’s effectiveness in notably diminishing this discrepancy. As shown in Figure 4, CLiTSSA ameliorates the performance differential between the French monolingual environment and cross-lingual context by 10.53, 3.31, and 2.13 absolute points in terms of F1-score for the L1, L2, and L3 tasks, respectively, thereby comparing the results to those observed in a monolingual context. Contrastingly, the English monolingual scenario exhibits a significant divergence from its French counterpart for L1 and L2 tasks, underscoring the imperative for speedy enhancements to bolster performance in monolingual contexts for languages with limited resources.

Error Analysis

Table 6 shows a couple of instances underscoring the challenges of semantic alignment. Initially, the heightened time-sensitive alignment offered by our model does not rectify the **inaccuracies in the foundational knowledge** of the underlying LLM. The first example elucidates that the factual inaccuracies inherent in the LLaMA3-8B model within a resource-rich linguistic context (i.e., English monolingual, En_m) persist despite the application of CLiTSSA in a French cross-lingual setting (Fr_c). Additionally, our

#	Setup	Question/Answer (Q/A)	Predicted Answer
1	X, En_m	Q: Which team did Glynn Snodin play for in Oct, 1986?, A: Sheffield Wednesday F.C.	Leeds United F.C.
	C, Fr_c	Q: Pour quelle équipe a joué Glynn Snodin en octobre 1986?, A: Sheffield Wednesday F.C.	Leeds United
2	X, Fr_m	Q: Qui était l'entraîneur principal de l'équipe HNK Rijeka après Simon Roman?, A: Goran Tomi	Ivan Pudar
	C, Fr_c	Q: Qui était l'entraîneur principal de l'équipe HNK Rijeka après Simon Roman?, A: Goran Tomi	Ivan Pudar

Table 6: Failure cases with CLiTSSA and their corresponding responses from X-InSTA in the monolingual scenario. En_m : English monolingual, Fr_m : French monolingual, and Fr_c : French cross-lingual. X: X-InSTA, and C: CLiTSSA




Figure 4: A comparative analysis of F1 scores across temporal tasks in monolingual and cross-lingual scenarios utilizing LLaMA3-8B, where En_m and Fr_m represent monolingual settings for English and French, respectively, while Fr_c is French’s cross-lingual setting.

proposed methodology is contingent on the semantic context within the monolingual embedding space, aligning the cross-lingual space accordingly. Consequently, **inaccuracies in expected responses may propagate** from the monolingual to the cross-lingual space notwithstanding the enhanced query alignment. A subsequent example illustrates this phenomenon in the contexts of French monolingual (Fr_m) and French cross-lingual scenarios. Furthermore, notwithstanding the semantic alignment ingrained in cross-lingual queries, the **implicit aspect of temporality** persistently presents a challenge, as observed for the L3 task.

Related Works

In NLP, significant foundational efforts in temporal reasoning encompass the creation of TimeBank (Pustejovsky et al. 2003), TempEval (Verhagen et al. 2010), and Time-stamped Language Models (Rajaby Faghihi and Kordjamshidi 2021), each contributing significantly to the understanding and processing temporal data. Concurrently, the evolution of knowledge graphs (KGs) has accentuated the importance of temporal relations therein, catalyzing the emergence of Temporal Knowledge Graph Completion (TKGC) as a distinct area of study. This progression has given rise to noteworthy question-answering datasets predicated on TKG, including TEQUILA (Jia et al. 2018), TimeQuestions (Jia et al. 2021), and CronQuestions (Saxena, Chakrabarti, and Talukdar 2021). The widespread use of language models in the public sphere further underscores the necessity for both temporal accuracy and consistency within generated responses. In response to this demand, several time-sensitive QA datasets, such as TEMPLAMA (Dhingra et al. 2022) and TEMPREASON (Tan, Ng, and Bing 2023), have been intro-

duced to assess and benchmark the temporal reasoning capabilities inherent in LLMs. Among these, TEMPREASON stands out as a comprehensive benchmark for temporal reasoning, spanning a broad spectrum of temporal periods and incorporating three levels of temporal relations. Further, the TEMP-COFAC (Bajpai et al. 2024) has been introduced to assess the temporally consistent factuality.

Furthermore, most LLMs are trained on multilingual datasets (Wenzek et al. 2020), a practice that was once a rarity given the dominance of extensive English corpora (Radford et al. 2019). Despite this, LLMs have proven their mettle in considerable languages. While there have been significant advancements in the multilingual capabilities of LLMs (Lin et al. 2022a; Qin et al. 2024), they still face significant challenges when dealing with low-resource languages (Cahyawijaya, Lovenia, and Fung 2024), especially in task-specific contexts (Enis and Hopkins 2024). To address this, innovative approaches such as prompting for generating intermediate English contexts (Huang et al. 2023), cross-lingual prompting (Winata et al. 2021), and Linguistically Diverse Prompting (LDP) (Nguyen et al. 2024) have been introduced. In the cross-lingual prompting domain, specific developments like semantic label-based alignment (Tanwar et al. 2023), query-based alignment via translation semantic similarity (Cahyawijaya, Lovenia, and Fung 2024), and a model-specific fine-tuned retriever (Lin, Martins, and Schütze 2024) have further enhanced LLMs’ capabilities. Yet, the exploration of temporal reasoning within low-resource languages remains scant, presenting a compelling area for further research. This study proposes to pioneer advancements in this under-explored domain.

Conclusion

In this paper, we introduced a novel dataset, m TEMPREASON, aimed at improving temporal reasoning assessment in low-resource languages using LLMs. Our analyses identified that multilingual LLMs inherently reward in-language time-sensitive semantic alignment over the cross-lingual similarity context in the X-ICL method. To overcome this, we proposed CLiTSSA, a novel method that enhances the retrieval of time-sensitive contextually relevant examples across low-resource languages. Our results demonstrated that this approach effectively improves LLMs’ temporal reasoning in low-resource languages, which we believe will aid in promoting linguistic diversity and the development of more inclusive LLMs. Future endeavors may benefit from examining the alignment between an implicit temporal query’s semantics and its implied semantic space to enhance intricate L3 task performance.

Acknowledgments

T. Chakraborty acknowledges the support of the IBM-IITD AI Horizons network and Rajiv Khemani Young Faculty Chair Professorship in Artificial Intelligence.

References

- Adilazuarda, M. F.; Cahyawijaya, S.; Aji, A. F.; Winata, G. I.; and Purwarianti, A. 2024. LinguAlchemy: Fusing Typological and Geographical Elements for Unseen Language Generalization. arXiv:2401.06034.
- AI@Meta. 2024. Llama 3 Model Card.
- Asai, A.; Kudugunta, S.; Yu, X. V.; Blevins, T.; Gonen, H.; Reid, M.; Tsvetkov, Y.; Ruder, S.; and Hajishirzi, H. 2023. BUFFET: Benchmarking Large Language Models for Few-shot Cross-lingual Transfer. arXiv:2305.14857.
- Bajpai, A.; Goyal, A.; Anwer, A.; and Chakraborty, T. 2024. Temporally Consistent Factuality Probing for Large Language Models. In AI-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 15864–15881. Miami, Florida, USA: Association for Computational Linguistics.
- Cahyawijaya, S.; Lovenia, H.; Aji, A. F.; Winata, G. I.; Willie, B.; Mahendra, R.; Wibisono, C.; Romadhony, A.; Vincentio, K.; Koto, F.; Santoso, J.; Moeljadi, D.; Wirawan, C.; Hudi, F.; Parmonangan, I. H.; Alfina, I.; Wicaksono, M. S.; Putra, I. F.; Rahmadani, S.; Oenang, Y.; Septiandri, A. A.; Jaya, J.; Dhole, K. D.; Suryani, A. A.; Putri, R. A.; Su, D.; Stevens, K.; Nityasya, M. N.; Adilazuarda, M. F.; Ignatius, R.; Diandaru, R.; Yu, T.; Ghifari, V.; Dai, W.; Xu, Y.; Damapuspita, D.; Tho, C.; Karo, I. M. K.; Fatyanosa, T. N.; Ji, Z.; Fung, P.; Neubig, G.; Baldwin, T.; Ruder, S.; Sujaini, H.; Sakti, S.; and Purwarianti, A. 2023. NusaCrowd: Open Source Initiative for Indonesian NLP Resources. arXiv:2212.09648.
- Cahyawijaya, S.; Lovenia, H.; and Fung, P. 2024. LLMs Are Few-Shot In-Context Low-Resource Language Learners. arXiv:2403.16512.
- Chen, W.; Wang, X.; and Wang, W. Y. 2021. A Dataset for Answering Time-Sensitive Questions. arXiv:2108.06314.
- Dhingra, B.; Cole, J. R.; Eisenschlos, J. M.; Gillick, D.; Eisenstein, J.; and Cohen, W. W. 2022. Time-Aware Language Models as Temporal Knowledge Bases. *Transactions of the Association for Computational Linguistics*, 10: 257–273.
- Enis, M.; and Hopkins, M. 2024. From LLM to NMT: Advancing Low-Resource Machine Translation with Claude. arXiv:2404.13813.
- Huang, H.; Tang, T.; Zhang, D.; Zhao, W. X.; Song, T.; Xia, Y.; and Wei, F. 2023. Not All Languages Are Created Equal in LLMs: Improving Multilingual Capability by Cross-Lingual-Thought Prompting. arXiv:2305.07004.
- Jia, Z.; Abujabal, A.; Saha Roy, R.; Strötgen, J.; and Weikum, G. 2018. TEQUILA: Temporal Question Answering over Knowledge Bases. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*. ACM.
- Jia, Z.; Pramanik, S.; Saha Roy, R.; and Weikum, G. 2021. Complex Temporal Question Answering on Knowledge Graphs. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management, CIKM '21*. ACM.
- Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; Lavaud, L. R.; Lachaux, M.-A.; Stock, P.; Scao, T. L.; Lavril, T.; Wang, T.; Lacroix, T.; and Sayed, W. E. 2023. Mistral 7B. arXiv:2310.06825.
- Lin, P.; Martins, A. F. T.; and Schütze, H. 2024. XAMPLER: Learning to Retrieve Cross-Lingual In-Context Examples. arXiv:2405.05116.
- Lin, X. V.; Mihaylov, T.; Artetxe, M.; Wang, T.; Chen, S.; Simig, D.; Ott, M.; Goyal, N.; Bhosale, S.; Du, J.; Pasunuru, R.; Shleifer, S.; Koura, P. S.; Chaudhary, V.; O’Horo, B.; Wang, J.; Zettlemoyer, L.; Kozareva, Z.; Diab, M.; Stoyanov, V.; and Li, X. 2022a. Few-shot Learning with Multilingual Generative Language Models. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 9019–9052. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Lin, X. V.; Mihaylov, T.; Artetxe, M.; Wang, T.; Chen, S.; Simig, D.; Ott, M.; Goyal, N.; Bhosale, S.; Du, J.; Pasunuru, R.; Shleifer, S.; Koura, P. S.; Chaudhary, V.; O’Horo, B.; Wang, J.; Zettlemoyer, L.; Kozareva, Z.; Diab, M.; Stoyanov, V.; and Li, X. 2022b. Few-shot Learning with Multilingual Language Models. arXiv:2112.10668.
- Liu, J.; Shen, D.; Zhang, Y.; Dolan, B.; Carin, L.; and Chen, W. 2022. What Makes Good In-Context Examples for GPT-3? In Agirre, E.; Apidianaki, M.; and Vulić, I., eds., *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, 100–114. Dublin, Ireland and Online: Association for Computational Linguistics.
- Miyabe, M.; and Yoshino, T. 2015. Evaluation of the Validity of Back-Translation as a Method of Assessing the Accuracy of Machine Translation. *2015 International Conference on Culture and Computing (Culture Computing)*, 145–150.
- Muennighoff, N.; Wang, T.; Sutawika, L.; Roberts, A.; Biderman, S.; Scao, T. L.; Bari, M. S.; Shen, S.; Yong, Z.-X.; Schoelkopf, H.; Tang, X.; Radev, D.; Aji, A. F.; Almubarak, K.; Albanie, S.; Alyafeai, Z.; Webson, A.; Raff, E.; and Raffel, C. 2023. Crosslingual Generalization through Multitask Finetuning. arXiv:2211.01786.
- Nguyen, X.-P.; Aljunied, S. M.; Joty, S.; and Bing, L. 2024. Democratizing LLMs for Low-Resource Languages by Leveraging their English Dominant Abilities with Linguistically-Diverse Prompts. arXiv:2306.11372.
- Pustejovsky, J.; Hanks, P.; Saurí, R.; See, A.; Gaizauskas, R.; Setzer, A.; Radev, D.; Sundheim, B.; Day, D.; Ferro, L.; and Lazo, M. 2003. The TimeBank corpus. *Proceedings of Corpus Linguistics*.
- Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; and Yu, P. S. 2024. Multilingual Large Lan-

- guage Model: A Survey of Resources, Taxonomy and Frontiers. arXiv:2404.04925.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language Models are Unsupervised Multitask Learners.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2023. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. arXiv:1910.10683.
- Rajaby Faghihi, H.; and Kordjamshidi, P. 2021. Time-Stamped Language Model: Teaching Language Models to Understand The Flow of Events. In Toutanova, K.; Rumshisky, A.; Zettlemoyer, L.; Hakkani-Tur, D.; Beltagy, I.; Bethard, S.; Cotterell, R.; Chakraborty, T.; and Zhou, Y., eds., *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4560–4570. Online: Association for Computational Linguistics.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. arXiv:1908.10084.
- Rubin, O.; Herzig, J.; and Berant, J. 2022. Learning To Retrieve Prompts for In-Context Learning. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2655–2671. Seattle, United States: Association for Computational Linguistics.
- Saxena, A.; Chakrabarti, S.; and Talukdar, P. 2021. Question Answering Over Temporal Knowledge Graphs. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 6663–6676. Online: Association for Computational Linguistics.
- Tan, Q.; Ng, H. T.; and Bing, L. 2023. Towards Benchmarking and Improving the Temporal Reasoning Capability of Large Language Models. arXiv:2306.08952.
- Tanwar, E.; Dutta, S.; Borthakur, M.; and Chakraborty, T. 2023. Multilingual LLMs are Better Cross-lingual In-context Learners with Alignment. arXiv:2305.05940.
- Verhagen, M.; Saurí, R.; Caselli, T.; and Pustejovsky, J. 2010. SemEval-2010 Task 13: TempEval-2. In Erk, K.; and Strapparava, C., eds., *Proceedings of the 5th International Workshop on Semantic Evaluation*, 57–62. Uppsala, Sweden: Association for Computational Linguistics.
- Wenzek, G.; Lachaux, M.-A.; Conneau, A.; Chaudhary, V.; Guzmán, F.; Joulin, A.; and Grave, E. 2020. CCNet: Extracting High Quality Monolingual Datasets from Web Crawl Data. In Calzolari, N.; Béchet, F.; Blache, P.; Choukri, K.; Cieri, C.; Declerck, T.; Goggi, S.; Isahara, H.; Maegaard, B.; Mariani, J.; Mazo, H.; Moreno, A.; Odijk, J.; and Piperidis, S., eds., *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 4003–4012. Marseille, France: European Language Resources Association. ISBN 979-10-95546-34-4.
- Winata, G. I.; Madotto, A.; Lin, Z.; Liu, R.; Yosinski, J.; and Fung, P. 2021. Language Models are Few-shot Multilingual Learners. In Ataman, D.; Birch, A.; Conneau, A.; Firat, O.; Ruder, S.; and Sahin, G. G., eds., *Proceedings of the 1st Workshop on Multilingual Representation Learning*, 1–15. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Yamada, I.; and Ri, R. 2024. LEIA: Facilitating Cross-lingual Knowledge Transfer in Language Models with Entity-based Data Augmentation. arXiv:2402.11485.
- Zhang, N.; Li, L.; Chen, X.; Deng, S.; Bi, Z.; Tan, C.; Huang, F.; and Chen, H. 2022. Differentiable Prompt Makes Pre-trained Language Models Better Few-shot Learners. arXiv:2108.13161.
- Zhao, T. Z.; Wallace, E.; Feng, S.; Klein, D.; and Singh, S. 2021. Calibrate Before Use: Improving Few-Shot Performance of Language Models. arXiv:2102.09690.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E. P.; Zhang, H.; Gonzalez, J. E.; and Stoica, I. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. arXiv:2306.05685.