

FrameEOL: Semantic Frame Induction using Causal Language Models

Chihiro Yano¹, Kosuke Yamada^{1,2}, Hayato Tsukagoshi¹,
Ryohei Sasano¹, Koichi Takeda³

¹Nagoya University ²CyberAgent ³National Institute of Informatics
yano.chihiro.j3@s.mail.nagoya-u.ac.jp, yamada_kosuke@cyberagent.co.jp
tsukagoshi.hayato.r2@s.mail.nagoya-u.ac.jp
sasano@i.nagoya-u.ac.jp takedasu@nii.ac.jp

Abstract

Semantic frame induction is the task of clustering frame-evoking words according to the semantic frames they evoke. In recent years, leveraging embeddings of frame-evoking words that are obtained using masked language models (MLMs) such as BERT has led to high-performance semantic frame induction. Although causal language models (CLMs) such as the GPT and Llama series succeed in a wide range of language comprehension tasks and can engage in dialogue as if they understood frames, they have not yet been applied to semantic frame induction. We propose a new method for semantic frame induction based on CLMs. Specifically, we introduce FrameEOL, a prompt-based method for obtaining Frame Embeddings that outputs **One** frame-name as a Label representing the given situation.¹ To obtain embeddings more suitable for frame induction, we leverage in-context learning (ICL) and deep metric learning (DML). Frame induction is then performed by clustering the resulting embeddings. Experimental results on the English and Japanese FrameNet datasets demonstrate that the proposed methods outperform existing frame induction methods. In particular, for Japanese, which lacks extensive frame resources, the CLM-based method using only 5 ICL examples achieved comparable performance to the MLM-based method fine-tuned with DML.

1 Introduction

Frame semantics (Fillmore, 1982) considers access to world knowledge that is related to a word to be essential for understanding the word. Consider the following sentences:

1. He *lost* the gold medal by just .02 points.
2. He *lost* his gold medal at the restaurant.

¹FrameEOL is a method inspired by PromptEOL, and its name includes the string “EOL” to reflect this connection. However, it should be noted that what “EOL” stands for differs between the two methods.

Both sentences contain the same verb *lost*. However, Sentence (1) conjures an image of “a competition comes to an end, with someone losing against an opponent according to their score,” while Sentence (2) evokes a situation where “someone may have unintentionally misplaced a possession.” When reading such sentences, humans recognize the different meanings of *lost* from the context and understand each sentence while supplying the information necessary to understand the situation, such as COMPETITION or OPPONENT for Sentence (1). In frame semantics, such background knowledge is referred to as a semantic frame and words like *lost* are called frame-evoking words. The verb *lost* in Sentence (1) corresponds to the FINISH_COMPETITION frame, which involves a competition coming to an end with an outcome determined by a score or ranking. In contrast, the verb *lost* in Sentence (2) corresponds to the LOSING frame, which describes a situation where someone loses his or her possession.

Semantic frame knowledge is being developed for many languages. The representative resource is FrameNet (Ruppenhofer et al., 2016), a large-scale semantic frame resource that was manually developed for English. In addition, resource development following FrameNet is in progress for several languages, including Brazilian Portuguese, Chinese, Hebrew, Japanese, Korean, and Swedish (Baker et al., 2018). However, manual resource development entails enormous costs, making it challenging to construct such large-scale frame knowledge. Accordingly, automatic construction of frame knowledge from massive text corpora is an active research area. Representative of such studies are those on frame induction: the task of clustering a set of words given together with their context according to the semantic frames they evoke. In recent years, several methods (Ribeiro et al., 2019; Arefyev et al., 2019; Yamada et al., 2021) that leverage masked language models

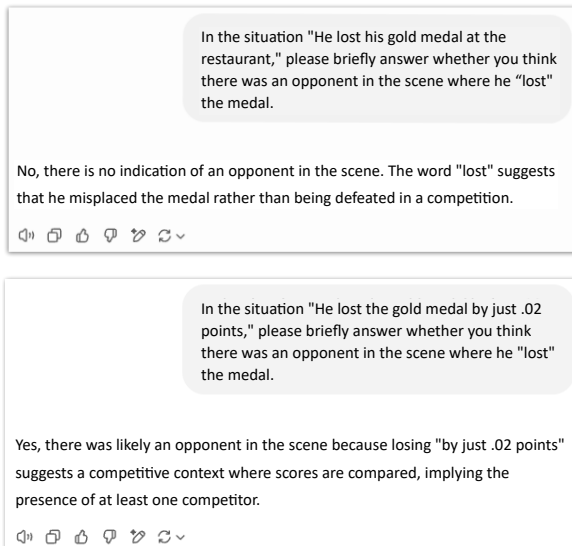


Figure 1: Examples of ChatGPT output showing the possibility of recognizing differences in semantic frames evoked by the same predicate *lost*.

(MLMs) such as BERT (Devlin et al., 2019) to obtain frame embeddings have been proposed. Frame embeddings are vector representations that capture the semantic frame evoked by a given word and are used to induce semantic frames.

In contrast, large-scale causal language models (CLMs) such as the GPT and Llama series have not been used for semantic frame induction, despite their success in a wide range of language comprehension tasks (Touvron et al., 2023; OpenAI, 2024). In fact, modern CLMs such as GPT-4o are potentially likely to understand semantic frames. For example, as shown in Figure 1, when ChatGPT receives the following prompt template with [sentence] replaced by Sentence (1) and Sentence (2), respectively, it tends to answer “Yes” to the former and “No” to the latter.²

Prompt template:

In the situation "[sentence]," please briefly answer whether you think there was an opponent in the scene where he "lost" the medal.

This indicates that GPT-4o correctly understands that an opponent is present only in the situation represented by Sentence (1). It is thus conceivable that frame induction could be performed with higher accuracy by leveraging such CLMs.

Hence, in this paper, we propose a method for inducing semantic frames based on CLMs. We first introduce **FrameEOL**, a prompt-based method for

obtaining **Frame Embeddings** that outputs **One** frame-name as a **Label** representing the given situation. FrameEOL extends PromptEOL (Jiang et al., 2024), a prompt-based sentence embedding method with an explicit one-word limitation, for semantic frame induction. Figure 2 shows an overview of FrameEOL and the overall frame induction process. Given a frame-evoking word and its context, FrameEOL inputs a prompt to CLMs such that the next word will be the corresponding frame name in FrameNet. As a result, the last token’s embedding is considered to be the embedding that would represent that frame. Furthermore, we apply in-context learning (ICL) and deep metric learning (DML) to obtain embeddings that are more suitable for semantic frame induction. After obtaining frame embeddings for the frame-evoking words in a given corpus, we cluster these embeddings to induce semantic frames.

2 Related Work

2.1 Semantic Frame Induction

Semantic frame induction is a task that groups frame-evoking words, typically verbs, along with their contexts according to the semantic frames they evoke.

Unlike semantic frame identification, which depends on a predefined set of frame labels, frame induction does not require such a predefined label set. This can be advantageous depending on the setting or application. First, when new concepts emerge, for example streaming in the sense of watching online videos, frame induction makes it possible to construct frames corresponding to those concepts (Tsujimoto et al., 2025). Second, in specific domains such as medicine or cooking, domain-specific concepts may exist and the required granularity of frames may differ. In such cases, it becomes necessary to perform frame induction that does not rely on a predefined label set.

In recent years, methods leveraging contextualized embeddings have been proposed and shown to achieve strong performance (Arefyev et al., 2019; Ribeiro et al., 2019; Anwar et al., 2019; Yamada et al., 2021, 2023; Mosolova et al., 2024). They leverage MLMs such as BERT (Devlin et al., 2019) to obtain contextualized word embeddings. The methods can be broadly categorized into two groups: unsupervised and supervised approaches. The unsupervised setting assumes that manually annotated frame information is unavailable, whereas

²We accessed <https://chatgpt.com/> on 2025-02-11 and used the GPT-4o model.

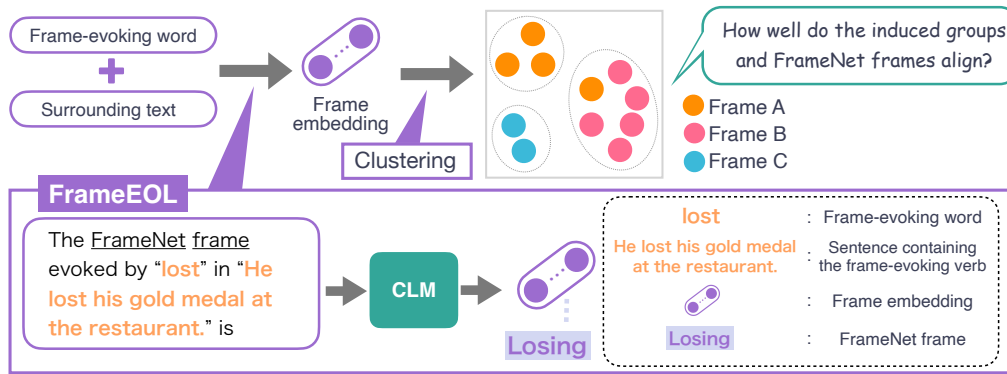


Figure 2: Overview of the frame induction process using FrameEOL and its evaluation.

the supervised setting assumes that manually annotated information is available for some frames.

In the unsupervised setting, several approaches have been proposed in recent years to group frame-evoking words via contextualized word embeddings derived from pre-trained language models (Arefyev et al., 2019; Ribeiro et al., 2019; Anwar et al., 2019). Among these, Yamada et al. (2021) utilized two types of embeddings from BERT: the contextualized word embeddings of frame-evoking words; and the embeddings when the frame-evoking words are masked with [MASK] tokens. Furthermore, to address the tendency of existing methods to overly disperse examples across too many clusters, they introduced a two-step clustering approach. While the commonly used one-step clustering approach clusters all examples at once, this approach first clusters examples by verb and subsequently clusters them across verbs, effectively limiting the number of clusters that examples of a single verb can belong to. In this study, we will conduct experiments with both the one-step clustering approach and the two-step clustering approach in semantic frame induction.

In the supervised setting, Yamada et al. (2023) demonstrated that semantic frame induction can be performed with even higher performance by applying deep metric learning (DML) (Kaya and Bilge, 2019). This approach addresses the limitation that embedding spaces obtained solely from pre-training do not always reflect meanings that align with human intuition about semantic frames. Through DML, the model is fine-tuned to place embeddings of verbs that evoke the same frame close together, while positioning embeddings of verbs evoking different frames farther apart in the semantic space. They employed several types of loss functions for DML including triplet loss, which

optimizes triplets of anchor, positive, and negative examples to ensure the distance between the anchor and negative example is at least a margin larger than the distance between the anchor and positive example.

2.2 Prompt-Based Text Embeddings

In this study, we adapt prompt-based embedding techniques originally developed for sentence representations, to derive dense embeddings that represent semantic frames. Therefore, in this section, we first provide an overview of existing methods for obtaining sentence embeddings.

Traditional methods typically construct sentence embeddings with various strategies, including extracting representations from special tokens (e.g., [CLS]) or computing the average of all token embeddings (Reimers and Gurevych, 2019; Tsukagoshi et al., 2021; Gao et al., 2021).

More recently, prompt-based approaches have been explored to enhance the embedding quality (Jiang et al., 2022, 2024; Lei et al., 2024). PromptBERT (Jiang et al., 2022) leverages a masked token prediction capability to obtain more effective sentence embeddings by using prompts such as “This sentence: [sentence] means [MASK]”, which consolidate the sentence’s semantic content into the [MASK] token representation. By leveraging masked token prediction, this approach dynamically contextualizes sentence meaning and has been shown to improve embedding quality.

PromptEOL (Jiang et al., 2024) is a method to obtain sentence embeddings by utilizing the next token prediction capability of CLMs, and it demonstrates high performance with optimization via ICL or contrastive learning. In this approach, a prompt is fed to a CLM to predict a single-word output

that represents a given sentence’s meaning. The last token’s final layer representation is used as the sentence embedding. The following is an example of the prompt used in PromptEOL.

Prompt template in PromptEOL:

This sentence: "[sentence]" means in one word: "

The placeholder [sentence] represents the sentence to be embedded.

Both approaches design the input format to closely resemble the pre-training task, thereby enabling the model to better understand the task and obtain more effective embeddings. Inspired by these methods, we design tailored prompts for CLMs to effectively obtain frame embeddings.

3 Proposed Method

We induce semantic frames by obtaining the frame embeddings of frame-evoking words with FrameEOL and then clustering the frame-evoking words in a corpus based on the obtained frame embeddings. Following Yamada et al. (2021), we use two clustering methods: one-step and two-step clustering. In this section, we first explain FrameEOL, a method for obtaining frame embeddings via a CLM; and then we describe how to optimize frame embeddings via ICL and DML.

3.1 FrameEOL

Inspired by PromptEOL (Jiang et al., 2024), we obtain the embedding of a frame-evoking word in a sentence by using a prompt tailored for frame induction. In PromptEOL, the prompt is designed to predict a single word representing the given sentence as the next token, whereas in FrameEOL, the prompt is designed to predict a FrameNet frame name as the next token. Specifically, FrameEOL uses the following prompt.

Prompt template in FrameEOL:

The FrameNet frame evoked by "[verb]" in "[sentence]" is

Here, [verb] denotes a placeholder for the frame-evoking verb for which we obtain the embedding, and [sentence] represents a placeholder for a sentence containing that frame-evoking verb. We use the embedding from the final layer corresponding to the last token “is” as the frame embedding.

Note that FrameEOL assumes that pre-trained CLMs have learned knowledge of FrameNet, as can be seen from the inclusion of the word “FrameNet”




Figure 3: Example of input in 3-shot ICL setting.

in the above prompt. This could be seen as some sort of leakage. In Section 4.3, we remove the explicit term “FrameNet” from the input prompt and show that, under both ICL and DML settings, omitting this term has only a limited effect on performance. Since this study focuses on frame induction in a supervised setting, we believe this potential leakage is not a major issue.

3.2 Frame Embedding Optimization via ICL

Semantic frame induction has primarily been studied in English. Its application to other languages remains challenging because of the limited availability of semantic frame resources. ICL is a method that adapts CLMs to downstream tasks by incorporating a few demonstrations into the input (Brown et al., 2020). It has gained significant attention, particularly in scenarios with limited training data.

To achieve high-performance frame induction even for low-resource languages, we introduce a method to obtain frame embeddings based on ICL. Figure 3 shows an example input provided to CLMs in a three-shot setting in our approach. We first prepare examples comprising a frame-evoking word, its surrounding sentence, and the evoked frame name. Next, we construct demonstrations by filling the FrameEOL prompt template: we fill the frame-evoking word into the [verb], fill the surrounding sentence into the [sentence], and then append the evoked frame name immediately afterward. Following these demonstrations, we append an inference prompt that is constructed by filling the FrameEOL template with the target instance’s frame-evoking word and its surrounding sentence without appending an evoked frame name. By including such demonstrations, we expect to output embeddings that represent frame names, thus yielding distinct embeddings for each evoked frame.

3.3 Frame Embedding Optimization via DML

Previous studies on semantic frame induction using MLMs have demonstrated that, when supervised data are available, fine-tuning through DML

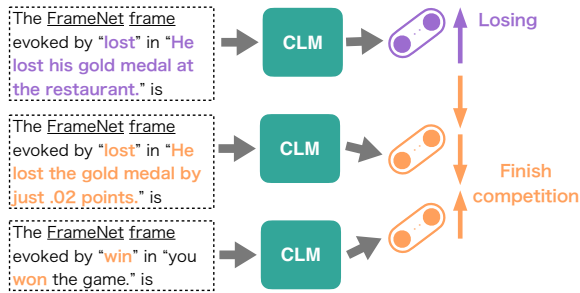


Figure 4: Overview of DML: embeddings of the same frame are trained to be closer, while embeddings of different frames are trained to be farther apart.

is more effective than using pre-trained models as they are (Yamada et al., 2023). Given its proven effectiveness in prior work, we expect DML to be similarly effective here. Hence, we use DML to enhance semantic frame induction using CLMs.

Given its demonstrated effectiveness (Yamada et al., 2023), we use the triplet loss:

$$\mathcal{L}_{\text{tri}} = \max(D(x_a, x_p) - D(x_a, x_n) + m, 0), \quad (1)$$

where x_a , x_p , and x_n are the frame embeddings of the anchor, positive, and negative instances, respectively. The margin m is a tunable hyperparameter, and $D(\cdot, \cdot)$ denotes a distance metric. In this study, we use the Euclidean distance of normalized frame embeddings x , following Yamada et al. (2023). Figure 4 shows an overview of our DML application.

Because we use large-scale CLMs, we adopt a parameter-efficient approach to fine-tune these models under a constraint of limited computational resources. Specifically, we use Low-Rank Adaptation (LoRA), a method that makes only the introduced low-rank matrices in linear layers trainable (Hu et al., 2022). LoRA has been shown to achieve performance comparable to that of full fine-tuning while significantly reducing the number of trainable parameters.

4 Experiments on FrameNet

We first conducted experiments on semantic frame induction using CLMs for English. In these experiments, we used FrameNet, an English frame knowledge resource, as the dataset for evaluation.

4.1 Evaluation Setup

Dataset We used a dataset that we extracted from FrameNet 1.7 (Ruppenhofer et al., 2016) by selecting in which the frame-evoking word is a verb. For three-fold cross validation, the dataset was equally split into three subsets in a 1:1:1 ratio for training,

development, and test. The splits ensured that all examples of the same verb were placed in the same subset, and that the proportion of polysemous verbs remained constant across all subsets. The statistics for the overall dataset and each subset are provided in Appendix A. The average number of instances across the three subsets was 27,537. A key characteristic of the experimental setup is that the test set includes frames that were unseen during training. For instance, in the FrameNet dataset, an average of 135.3 out of 434.3 frames per test set were unseen during training.

Clustering Methods In the semantic frame induction task, the instances of each verb are clustered by the semantic frame that it evokes. Following previous work (Yamada et al., 2023, 2021), we employ two clustering methods in our experiments: one-step clustering and two-step clustering. For one-step clustering, we used group-average clustering, whereas for two-step clustering, we used X-means (Pelleg and Moore, 2000) for the first step and group-average clustering for the second step. Details are shown in Appendix B.

Evaluation Methods We evaluated the semantic frame induction methods by comparing automatically created clusters of instances of verbs with gold clusters of instances of verbs evoking the same frames. For evaluation metrics, we used B-cubed Precision (BCP), Recall (BCR), and their harmonic mean, F-score (BCF) (Bagga and Baldwin, 1998), which are representative metrics for this clustering task. In our experiments, we use BCF as the primary evaluation metric.

Training Settings In our experiments, we used widely recognized, general-purpose pre-trained language models from the Gemma and Llama series.

Specifically, we conducted experiments with Gemma 3-12B³ (Gemma Team, 2025), and Llama 3.1-8B⁴ (Llama Team, 2024). Experiments on smaller models are presented in Appendix C, due to space constraints. We conducted fine-tuning with LoRA by attaching LoRA adapters to all linear layers within the model. We set LoRA rank $r = 8$ and $\alpha = 32$, and training for 20 epochs via the AdamW optimizer (Loshchilov and Hutter, 2017) with a batch size of 32. We explored triplet loss

³<https://huggingface.co/google/gemma-3-12b-it>

⁴<https://huggingface.co/meta-llama/Llama-3.1-8B>

Model	Model Size	Training		One-Step Clustering			Two-Step Clustering		
		Method	Data Size	BcP	BcR	BcF	BcP	BcR	BcF
MLM-based methods									
BERT _{base}	0.1B	-	-	41.1	44.7	42.8	49.7	65.8	56.5
		DML	27,536	61.8	68.7	65.0	65.6	67.8	66.4
BERT _{large}	0.3B	-	-	32.8	45.9	38.3	51.2	65.4	57.3
		DML	27,536	63.6	69.0	66.2	67.3	69.2	67.8
ModernBERT _{base}	0.1B	-	-	19.6	42.8	27.0	50.7	61.4	55.4
		DML	27,536	59.7	69.7	64.3	65.1	68.4	66.5
ModernBERT _{large}	0.3B	-	-	29.4	45.1	35.6	51.6	62.1	56.2
		DML	27,536	63.5	73.1	67.9	68.2	68.7	68.2
RoBERTa _{large}	0.3B	-	-	21.9	58.0	31.7	52.9	63.8	57.6
		DML	27,536	62.9	73.7	67.9	70.4	69.2	69.6
CLM-based methods (ours)									
Gemma 3	12B	ICL	0	9.1	66.7	16.0	53.3	39.6	45.4
			5	62.2	55.4	58.6	64.0	53.8	58.4
			10	63.0	54.9	58.7	64.5	53.7	58.6
			20	63.0	56.7	59.7	64.8	54.5	59.2
			DML	27,536	69.3	74.7	71.9	72.2	69.4
Llama 3.1	8B	ICL	0	20.8	45.1	28.4	43.4	48.8	46.0
			5	55.9	58.3	57.0	49.0	62.8	54.9
			10	56.6	59.0	57.7	50.9	62.9	56.1
			20	56.8	60.1	58.3	51.5	63.3	56.7
			DML	27,536	67.2	74.9	70.8	74.2	68.3

Table 1: The experimental results of semantic frame induction on FrameNet. The scores were averaged over a three-fold cross-validation, and then multiplied by 100.

margins of {0.1, 0.2, 0.5, 1.0} and learning rates of {3e-5, 5e-5, 1e-4}. All hyperparameters were determined using the development set. Training cost details are shown in Appendix D.

In the ICL experiments, the demonstrations included in the prompt were randomly sampled from the training data. We experimented with 5, 10, and 20 demonstrations, conducting four runs with different random seeds to account for sampling variance. We set the maximum sequence length to 2048 and truncated the beginnings of input sequences when necessary.⁵ To prevent truncation, we replaced any demonstration exceeding 1,900 tokens with another randomly selected example.

Comparison Methods For comparison, we evaluated the method leveraging mask embeddings to obtain contextualized word embeddings (Yamada et al., 2021, 2023), and also varied the choice of pre-trained models, including BERT_{base}, BERT_{large}, ModernBERT_{base}, ModernBERT_{large}, and RoBERTa_{large}. ModernBERT is a powerful MLM that incorporates optimization techniques recently applied to CLMs (Warner et al., 2025).

We evaluated methods without fine-tuning and methods with fine-tuning by DML using the triplet loss.

⁵According to the Llama tokenizer, an average token length of 26.48 per example, and the prompt token length was 13

4.2 Results

Table 1 summarizes the results. Our proposed FrameEOL-based method achieved the highest performance among all approaches by applying DML to the model. Specifically, in the one-step clustering setting, Gemma 3 achieved a BcF score of 71.9 and Llama 3.1 scored 70.8—both substantially higher than the best encoder model’s score of 67.9. In the two-step clustering setting, Gemma 3 and Llama 3.1 attained BcF scores of 70.6 and 70.9, respectively, again surpassing the MLM-based method’s score of 69.6.

The performance of the methods using ICL tended to improve slightly as the number of demonstrations increased. The results confirmed that the CLM-based methods performed worse than the MLM-based methods when training data were not used, but the CLM-based methods achieved slightly better performance than the MLM-based methods when at least 20 instances were prepared. Specifically, the MLM-based method attained a score of 57.6 in the no-tuning setting, whereas the CLM-based method attained scores of 59.7 and 58.3 in the 20-shot ICL setting. The ICL-based methods exhibit particular potential for use with low-resource language.

When using DML or ICL, no performance improvement was observed with two-step clustering. This suggests that the learning process optimized

Model	Training Setting	“FrameNet” in input	One-Step BcF	Two-Step BcF
Gemma 3 -12B	-	Yes	16.0	45.4
		No	13.5 (-1.5)	43.9 (-1.5)
	ICL 20-shot	Yes	59.7	59.2
		No	59.6 (-0.1)	59.5 (+0.3)
	DML	Yes	71.9	70.6
		No	71.2 (-0.7)	70.8 (+0.2)
Llama 3 -8B	-	Yes	28.4	46.0
		No	8.8 (-19.6)	35.9 (-10.1)
	ICL 20-shot	Yes	58.3	56.7
		No	58.2 (-0.1)	55.8 (-0.9)
	DML	Yes	70.8	70.9
		No	70.5 (-0.3)	70.4 (-0.5)

Table 2: Differences in performance when the term “FrameNet” was excluded from the input.

the semantic embedding space sufficiently for semantic frames, thus enabling effective clustering of verb instances even with the one-step approach.

4.3 Effect of “FrameNet” in the Prompt

FrameEOL explicitly incorporates FrameNet semantics by instructing the prompt to generate FrameNet frames for subsequent tokens in the prompt. Given that the CLM’s pre-training resources likely include FrameNet, the model’s ability to perform reasonably well on few-shot semantic frame induction can be attributed to its potential use of FrameNet-related knowledge acquired during the pre-training phase. In this section, we examine the performance impact of the following input template, which omits the term “FrameNet” from the FrameEOL template.

Prompt template without “FrameNet”:

The frame evoked by “[verb]” in “[sentence]” is

If the model can achieve comparable performance without the explicit term “FrameNet”, then we can rule out the possibility that it is solving the task by leaking FrameNet information acquired during pre-training.

Table 2 lists the results. In the setting without fine-tuning, the effect on performance of removing the term “FrameNet” from the input prompt varies by model; however, under both the ICL and DML settings, this effect remains limited. These findings demonstrate that FrameNet information is not directly leaked from the pre-training corpus through the use of the term “FrameNet”. In other words, even if the pre-training corpus contains no explicit FrameNet annotations or related content, a CLM can still achieve high performance in frame induc-

tion once it is provided with sufficient annotated training data.

5 Experiments on Japanese FrameNet

For languages besides English, the limited scale of frame knowledge resources has constrained the development of semantic frame induction methods. Consequently, most research in this field has been conducted in English. For instance, Japanese FrameNet (Ohara et al., 2004), a frame knowledge resource for Japanese, contains only 3,130 examples where verbs serve as frame-evoking words. This is a very small dataset, representing only about 3.2% of the number of similar examples in FrameNet. The scarcity of training data makes it challenging to effectively train deep learning models, which has hindered the advancement of semantic frame induction methods in non-English languages.

In this study, to explore the performance of semantic frame induction for non-English languages, we also performed experiments using Japanese FrameNet for training and evaluation. Large-scale CLMs have demonstrated strong multilingual capabilities, even when they are primarily English-centric (Brown et al., 2020; Holtermann et al., 2024). Accordingly, we investigated whether leveraging these multilingual capabilities can be beneficial for semantic frame induction, even when the amount of training data is limited.

5.1 Experimental Setup

Our proposed method, FrameEOL, is adaptable to any language by simply modifying the prompt. To obtain frame embeddings in Japanese, we used the following prompt, which is a Japanese translation of the prompt described in Section 3.1.

Prompt template for Japanese:

“[sentence]” 内の “[verb]” が喚起する FrameNet フレームは

(English Translation: The FrameNet frame evoked by “[verb]” in “[sentence]” is)

Experiments were conducted with two general-purpose CLMs, Gemma 3-12B-it (Gemma Team, 2025) and Llama 3.1-8B (Llama Team, 2024), and one CLM intended for use in Japanese, LLM-jp-3-13B-Instruct3⁶ (LLM-jp, 2024). LLM-jp-3-13B-Instruct3 was trained on a total of 2.1T tokens,

⁶<https://huggingface.co/llm-jp/llm-jp-3-13b-instruct3>

Model	Model Size	Training		One-Step Clustering			Two-Step Clustering		
		Method	Data Size	BcP	BcR	BcF	BcP	BcR	BcF
MLM-based methods									
Japanese BERT _{base}	0.1B	-	-	42.8	47.7	44.6	47.3	57.9	51.8
		DML	1,100	55.1	60.2	57.1	54.3	63.1	58.2
Japanese ModernBERT _{base}	0.1B	-	-	38.8	42.7	40.5	53.7	54.5	54.0
		DML	1,100	59.0	61.3	60.0	51.6	63.0	56.6
Japanese ModernBERT _{large}	0.1B	-	-	36.8	40.0	38.2	50.6	55.5	52.7
		DML	1,100	56.8	63.5	59.8	54.6	63.1	58.4
CLM-based methods (ours)									
Gemma 3	12B	ICL	0	16.9	21.9	18.7	39.9	45.0	42.3
			5	55.6	61.9	58.5	57.4	60.9	59.0
			10	56.4	61.0	58.5	57.8	61.1	59.4
			20	56.7	60.2	58.3	57.9	60.8	59.3
			DML	1,100	51.9	66.2	57.9	53.8	66.2
Llama 3.1	8B	ICL	0	18.9	43.7	26.2	47.2	50.2	48.5
			5	55.7	65.1	59.9	51.2	65.5	57.4
			10	56.9	64.9	60.6	51.8	65.6	57.8
			20	56.8	65.5	60.9	51.6	65.8	57.8
			DML	1,100	57.0	64.1	60.3	56.3	65.3
LLM-jp-3	13B	ICL	0	19.7	32.7	24.5	43.1	55.1	47.9
			5	53.1	65.0	58.4	53.0	64.6	58.2
			10	54.2	65.8	59.4	51.9	66.2	58.1
			20	53.8	66.3	59.3	51.7	67.1	58.4
			DML	1,100	56.9	66.7	61.3	53.4	66.6

Table 3: The Experimental results of semantic frame induction on Japanese FrameNet. The scores were averaged over a three-fold cross-validation, and then multiplied by 100.

which included Japanese data, and underwent instruction tuning and direct preference optimization (Rafael et al., 2023). The experiments on smaller models are presented in Appendix C. The other experimental settings for clustering methods, DML, and ICL were the same as those described in Section 4.1. The datasets and comparison methods used for evaluation were as follows.

Dataset We extracted Japanese FrameNet examples in which verbs evoked specific frames. For three-fold cross-validation, these examples were divided into training, development, and test sets, ensuring no overlap of frames or evoking words across subsets. For frames with three or more evoking words, overlaps between subsets were allowed to prevent the number of evoking words per frame from becoming too small, given the dataset’s limited size. Nevertheless, as the split was performed at the verb level, no single verb appeared in multiple subsets. The statistics for the overall dataset and each subset are provided in Appendix A. The average number of instances across the three subsets was 1,043. In the Japanese FrameNet dataset, an average of 110.7 out of 137.0 frames per test set were unseen during training.

Comparison Methods For comparison methods, we also evaluated MLM-based approaches that

work without fine-tuning (Yamada et al., 2021) and with fine-tuning (Yamada et al., 2023) to a Japanese BERT model and Japanese ModernBERT model. The BERT model was the one released by Tohoku University,⁷ which has 110M parameters and was trained on the Japanese Wikipedia and CC-100 datasets. This model is widely used in Japanese natural language processing tasks. The ModernBERT model,⁸ was trained with 4.39T tokens of Japanese and English data (Tsukagoshi et al., 2025), which is more than double the amount used for LLM-jp-3-13B-Instruct3.

5.2 Experimental Results

Table 3 summarizes the experimental results. Our proposed CLM-based methods demonstrated superior performance compared to the MLM-based methods, regardless of the training setting. The highest overall performance was achieved in the one-step clustering setting using LLMjp-3, reaching a BcF score of 61.3, surpassing the MLM-based method’s score of 60.0. In the two-step clustering setting, Llama 3.1 achieved the highest BcF score of 60.3, exceeding the MLM-based method’s score of 58.4. Across both clustering approaches, the

⁷<https://huggingface.co/cl-tohoku/bert-base-japanese-v3>

⁸<https://huggingface.co/sbintuitions/modernbert-ja-130m>

ICL setting using five or more few-shot examples demonstrated performance that was comparable to the method of fine-tuning MLM with DML. This was likely due to the small data size, as the training data for Japanese FrameNet comprises only 1,100 examples, approximately 4% of the 27,536 examples available in English FrameNet.

When we use ICL, sufficient performance can be achieved with as few as 20 examples, which is significantly smaller compared to the number of examples required for DML with MLM. Hence, even in cases where only a very limited number of examples is available, these results demonstrate the possibility of achieving highly accurate automatic construction of semantic frame knowledge by creating a small number of examples and leveraging CLMs combined with ICL.

6 Conclusion

In this study, we proposed a method using causal language models (CLMs) for semantic frame induction of verbs. By leveraging the characteristics of CLMs, our proposed method uses a sequence in which the next word is the frame name as input to the CLMs, and it uses the final layer embedding of the last token as the frame embedding representation. In experiments on FrameNet, deep metric learning (DML) exceeded the performance of existing methods using masked language models (MLMs) with DML. In-context learning (ICL) showed lower performance than existing methods using MLMs trained on a large amount of supervised data with DML, but better performance than existing methods without DML. In experiments on Japanese FrameNet, the ICL method demonstrated performance comparable to methods based on MLMs fine-tuned with DML. These results suggest that, when automatically constructing semantic frame knowledge for languages with no or very limited training data, a promising approach would be to create about 20 training examples and use a method combining CLMs with ICL.

Limitations

Our proposed CLM-based semantic frame induction method demonstrated high performance. Nevertheless, several limitations remain. In particular, our experiments were limited to English FrameNet and Japanese FrameNet, leaving it unclear how well our approach would generalize to other languages or corpora. Additionally, the use of large

CLMs entails significant computational requirements, which could restrict the practical deployment of our method in resource-constrained environments.

Acknowledgments

We would like to express our gratitude to Dr. Kyoko Ohara of Keio University for providing the Japanese FrameNet data used in this study. This research was supported by JST FOREST Program JPMJFR216N.

References

- Saba Anwar, Dmitry Ustalov, Nikolay Arefyev, Simone Paolo Ponzetto, Chris Biemann, and Alexander Panchenko. 2019. [HHMM at SemEval-2019 Task 2: Unsupervised Frame Induction using Contextualized Word Embeddings](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval 2019)*, pages 125–129.
- Nikolay Arefyev, Boris Sheludko, Adis Davletov, Dmitry Kharchev, Alex Nevidomsky, and Alexander Panchenko. 2019. [Neural GRANNy at SemEval-2019 Task 2: A Combined Approach for Better Modeling of Semantic Relationships in Semantic Frame Induction](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval 2019)*, pages 31–38.
- Amit Bagga and Breck Baldwin. 1998. [Entity-Based Cross-Document Coreferencing Using the Vector Space Model](#). In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL-COLING 1998)*, pages 79–85.
- Collin F. Baker, Michael Ellsworth, Miriam R. L. Petruck, and Swabha Swayamdipta. 2018. [Frame Semantics across Languages: Towards a Multilingual FrameNet](#). In *Proceedings of the 27th International Conference on Computational Linguistics: Tutorial Abstracts (COLING 2018)*, pages 9–12.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems (NeurIPS 2020)*, volume 33, pages 1877–1901.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of](#)

- Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 4171–4186.
- Charles J Fillmore. 1982. *Frame Semantics*. In *Linguistics in the Morning Calm*, pages 111–137. Hanshin Publishing Company.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. *SimCSE: Simple Contrastive Learning of Sentence Embeddings*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP 2021)*, pages 6894–6910.
- Gemma Team. 2025. *Gemma 3 Technical Report*. *arXiv preprint arXiv:2503.19786*.
- Carolin Holtermann, Paul Röttger, Timm Dill, and Anne Lauscher. 2024. *Evaluating the Elementary Multilingual Capabilities of Large Language Models with MultiQ*. In *Findings of the Association for Computational Linguistics (ACL 2024)*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. *LoRA: Low-Rank Adaptation of Large Language Models*. In *International Conference on Learning Representations (ICLR 2022)*.
- Ting Jiang, Shaohan Huang, Zhongzhi Luan, Deqing Wang, and Fuzhen Zhuang. 2024. *Scaling Sentence Embeddings with Large Language Models*. In *Findings of the Association for Computational Linguistics: EMNLP 2024*.
- Ting Jiang, Jian Jiao, Shaohan Huang, Zihan Zhang, Deqing Wang, Fuzhen Zhuang, Furu Wei, Haizhen Huang, Denvy Deng, and Qi Zhang. 2022. *PromptBERT: Improving BERT Sentence Embeddings with Prompts*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022)*.
- Mahmut Kaya and Hasan Şakir Bilge. 2019. *Deep Metric Learning: A Survey*. *Symmetry*, 11(9):1066.
- Yibin Lei, Di Wu, Tianyi Zhou, Tao Shen, Yu Cao, Chongyang Tao, and Andrew Yates. 2024. *Meta-Task Prompting Elicits Embeddings from Large Language Models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*.
- Llama Team. 2024. *The Llama 3 Herd of Models*. *arXiv preprint arXiv:2407.21783*.
- LLM-jp. 2024. *LLM-jp: A Cross-organizational Project for the Research and Development of Fully Open Japanese LLMs*. *arXiv preprint arXiv:2407.03963*.
- Ilya Loshchilov and Frank Hutter. 2017. *Decoupled Weight Decay Regularization*. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*.
- Anna Mosolova, Marie Candito, and Carlos Ramisch. 2024. *Injecting Wiktionary to improve token-level contextual representations using contrastive learning*. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024)*.
- Kyoko Ohara, Seiko Yamaguchi Fujii, Toshio Ohori, Ryoko Suzuki, Hiroaki Saito, and Shun Ishizaki. 2004. *The Japanese FrameNet Project: An Introduction*.
- OpenAI. 2024. *GPT-4 Technical Report*. *arXiv preprint arXiv:2408.00118*.
- Dan Pelleg and Andrew Moore. 2000. *X-means: Extending K-means with Efficient Estimation of the Number of Clusters*. In *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pages 727–734.
- Rafailov Rafael, Sharma Archit, Mitchell Eric, D Manning Christopher, Ermon Stefano, and Finn Chelsea. 2023. *Direct Preference Optimization: Your Language Model is Secretly a Reward Model*. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS 2023)*.
- Nils Reimers and Iryna Gurevych. 2019. *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019)*.
- Eugénio Ribeiro, Vânia Mendonça, Ricardo Ribeiro, David Martins de Matos, Alberto Sardinha, Ana Lúcia Santos, and Luísa Coheur. 2019. *L2F/INESC-ID at SemEval-2019 Task 2: Unsupervised Lexical Semantic Frame Induction using Contextualized Word Representations*. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval 2019)*, pages 130–136.
- Josef Ruppenhofer, Michael Ellsworth, Myriam Schwarzer-Petruck, Christopher R Johnson, and Jan Scheffczyk. 2016. *FrameNet II: Extended theory and practice*. International Computer Science Institute.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. *Llama 2: Open Foundation and Fine-Tuned Chat Models*. *arXiv preprint arXiv:2307.09288*.
- Shogo Tsujimoto, Kosuke Yamada, and Ryohei Sasano. 2025. *Semantic Frame Induction from a Real-World Corpus*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop) (ACL-SRW 2025)*.

Hayato Tsukagoshi, Shengzhe Li, Akihiko Fukuchi, and Tomohide Shibata. 2025. [ModernBERT-Ja](https://huggingface.co/collections/sbintuitions/modernbert-ja-67b68fe891132877cf67aa0a). <https://huggingface.co/collections/sbintuitions/modernbert-ja-67b68fe891132877cf67aa0a>.

Hayato Tsukagoshi, Ryohei Sasano, and Koichi Takeda. 2021. *DefSent: Sentence Embeddings using Definition Sentences*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Griffin Thomas Adams, Jeremy Howard, and Iacopo Poli. 2025. *Smarter, Better, Faster, Longer: A Modern Bidirectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference*. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL2025)*.

Kosuke Yamada, Ryohei Sasano, and Koichi Takeda. 2021. *Semantic Frame Induction using Masked Word Embeddings and Two-Step Clustering*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*, pages 811–816.

Kosuke Yamada, Ryohei Sasano, and Koichi Takeda. 2023. *Semantic Frame Induction with Deep Metric Learning*. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (ACL 2023)*, pages 1833–1845.

A Statistics of the datasets

	#Instances	#Frames	#Verbs
Split ₀	28,117	459	830
Split ₁	28,314	429	831
Split ₂	26,179	415	831
Full	82,610	642	2,492

Table 4: Statistics of the datasets used in the FrameNet experiments.

	#Instances	#Frames	#Verbs
Split ₀	1,046	112	209
Split ₁	1,055	154	294
Split ₂	1,029	145	263
Full	3,130	344	766

Table 5: Statistics of the datasets used in the Japanese FrameNet experiments.

The statistics of the datasets used in the experiments are presented. Table 4 shows the statistics for FrameNet, and Table 5 shows the statistics for

Japanese FrameNet. The average proportion of polysemous verbs in the test sets, which can evoke different frames depending on the context, was 30.3% for FrameNet and 19.1% for Japanese FrameNet.

B Clustering Setting

Following Yamada et al. (2021), we determine all clustering thresholds on the development set.

One-step clustering We employ group-average agglomerative clustering and iteratively merge clusters until their total number equals the number of frames in the development set. We then set the threshold to the minimum inter-cluster distance observed at this final merge.

Two-step clustering We begin by applying X-Means clustering to each group of examples sharing the same frame-evoking word, automatically inferring the optimal cluster count while capping the maximum clusters k_{\max} to the highest number of unique frames per verb in the development set. We then perform group-average agglomerative merging on the resulting clusters until the proportion of same-lemma examples assigned to the same cluster matches the proportion recorded on the development set. The minimum inter-cluster distance at this matching merge point serves as the second-stage threshold.

C Evaluation on Smaller Models

We report the frame induction performance of our proposed method when applied to smaller variants of the models used in Sections 4 and 5: Gemma 3-1B⁹, Gemma 3-4B¹⁰, Llama 3.2-1B¹¹, and Llama 3.2-3B¹². Table 6 shows the results on FrameNet, and Table 7 shows the results on Japanese FrameNet.

⁹<https://huggingface.co/google/gemma-3-1b-it>

¹⁰<https://huggingface.co/google/gemma-3-4b-it>

¹¹<https://huggingface.co/meta-llama/Llama-3.2-1B>

¹²<https://huggingface.co/meta-llama/Llama-3.2-3B>

Model	Size	Training		One-Step			Two-Step			
		Method	Data Size	BcP	BcR	BcF	BcP	BcR	BcF	
Gemma 3	1B	ICL	0	3.0	88.8	5.8	28.6	28.2	28.3	
			5	38.3	53.1	44.3	48.5	55.6	51.7	
			10	43.7	54.2	48.3	47.5	56.9	51.7	
			20	45.2	54.1	49.1	48.7	57.8	52.8	
	4B	DML	27,536	63.0	67.8	65.3	65.6	63.1	64.2	
			0	10.2	68.7	17.8	48.9	49.2	49.0	
			5	57.2	54.4	55.7	53.3	57.0	54.8	
			10	58.3	56.3	57.2	55.0	58.0	56.3	
4B	DML	20	58.1	59.3	58.6	56.9	59.5	57.9		
		27,536	67.7	68.8	68.2	69.2	67.3	68.1		
		0	3.0	84.4	5.9	46.0	30.2	36.4		
		5	24.9	57.4	34.1	44.2	59.4	50.6		
Llama 3.2	1B	ICL	10	30.2	59.0	39.7	44.4	59.9	50.9	
			20	34.6	59.8	43.7	44.6	60.8	51.4	
			DML	27,536	64.7	70.9	67.6	70.3	65.2	67.4
				0	12.6	69.1	21.3	54.8	45.1	49.5
	3B	ICL	5	44.3	59.4	50.6	47.7	64.2	54.7	
			10	47.3	59.4	52.5	48.4	62.9	54.6	
			20	48.8	60.3	53.9	48.2	63.3	54.7	
			DML	27,536	66.1	73.2	69.4	74.1	65.5	69.2

Table 6: Frame induction performance of smaller models on FrameNet. The scores were averaged over a three-fold cross-validation, and then multiplied by 100.

Model	Size	Training		One-Step			Two-Step			
		Method	Data Size	BcP	BcR	BcF	BcP	BcR	BcF	
Gemma 3	1B	ICL	0	16.8	19.2	17.8	37.9	44.0	40.7	
			5	45.8	60.8	52.1	50.8	60.0	54.9	
			10	46.8	61.8	53.1	49.1	61.7	54.6	
			20	49.0	62.4	54.8	48.9	62.9	54.9	
	4B	DML	1,100	54.1	62.5	57.8	52.4	63.9	57.5	
			0	16.5	20.4	17.9	40.6	47.2	43.6	
			5	52.4	65.5	58.2	55.1	63.3	58.8	
			10	52.5	66.0	58.4	56.2	63.5	59.5	
4B	DML	20	53.6	66.6	59.2	55.5	64.0	59.4		
		1,100	55.4	64.5	59.5	52.5	65.6	58.3		
		0	3.0	84.4	5.9	46.0	30.2	36.4		
		5	44.2	55.4	49.0	44.1	63.5	52.1		
Llama 3.2	1B	ICL	10	48.7	57.8	52.7	44.8	63.2	52.3	
			20	51.8	61.6	56.2	46.7	64.0	53.9	
			DML	1,100	55.1	63.2	58.8	52.9	64.2	58.0
				0	19.8	39.7	26.2	44.2	54.5	48.5
	3B	ICL	5	53.1	64.2	58.1	51.0	64.8	57.0	
			10	52.7	63.5	57.6	50.5	64.5	56.6	
			20	53.1	63.7	57.9	51.5	64.2	57.0	
			DML	1,100	55.5	65.7	60.1	55.9	64.9	60.0

Table 7: Frame induction performance of smaller models on Japanese FrameNet. The scores were averaged over a three-fold cross-validation, and then multiplied by 100.

D Training Cost

Model	Size	#Params	
		All	Trainable
Gemma 3	1B	1,006,408,832	6,522,880
	4B	3,895,164,416	14,901,248
	12B	11,798,769,408	32,735,232
Llama 3.2	1B	1,241,450,496	5,636,096
	3B	3,224,906,752	12,156,928
Llama 3.1	8B	7,525,896,192	20,971,520
LLM-jp-3	13B	13,229,347,840	31,293,440

Table 8: Total number of parameters and trainable parameters after adding the LoRA layers.

Dataset	Model	Size	Training Time
FrameNet	Gemma 3	1B	180 min
		4B	540 min
		12B	1200 min
	Llama 3.2	1B	150 min
		3B	360 min
	Llama 3.1	8B	720 min
Japanese FrameNet	Gemma 3	1B	12 min
		4B	30 min
		12B	80 min
	Llama 3.2	1B	12 min
		3B	30 min
	Llama 3.1	8B	60 min
LLM-jp-3	13B	60 min	

Table 9: Training time required for the DML experiments.

In deep metric learning, LoRA weights were applied to all linear layers. Table 8 shows the total and trainable parameters for each model after adding the LoRA layers.

The training was conducted on a single A100 GPU with 80GB VRAM, and the time required is shown in Table 9. For English FrameNet, an average of 27,537 instances were used for training; with a batch size of 32 over 20 epochs, the average number of training steps was 17,210. For Japanese FrameNet, an average of 1,043 instances were used for training; with a batch size of 32 over 20 epochs, the average number of training steps was 652.