

Few-Shot, No Problem: Descriptive Continual Relation Extraction

Nguyen Xuan Thanh^{1*}, Anh Duc Le^{2*}, Quyen Tran^{3*}, Thanh-Thien Le^{3*},
Linh Ngo Van^{2†}, Thien Huu Nguyen⁴

¹Oraichain Labs,

²Hanoi University of Science and Technology,

³VinAI Research,

⁴University of Oregon

thanh.nx@orai.io

anh.ld204628@sis.hust.edu.vn

{v.quyentt15,v.thientt3}@vinai.io,

linhnhv@soict.hust.edu.vn, thien@cs.uoregon.edu

Abstract

Few-shot Continual Relation Extraction is a crucial challenge for enabling AI systems to identify and adapt to evolving relationships in dynamic real-world domains. Traditional memory-based approaches often overfit to limited samples, failing to reinforce old knowledge, with the scarcity of data in few-shot scenarios further exacerbating these issues by hindering effective data augmentation in the latent space. In this paper, we propose a novel retrieval-based solution, starting with a large language model to generate descriptions for each relation. From these descriptions, we introduce a bi-encoder retrieval training paradigm to enrich both sample and class representation learning. Leveraging these enhanced representations, we design a retrieval-based prediction method where each sample "retrieves" the best fitting relation via a reciprocal rank fusion score that integrates both relation description vectors and class prototypes. Extensive experiments on multiple datasets demonstrate that our method significantly advances the state-of-the-art by maintaining robust performance across sequential tasks, effectively addressing catastrophic forgetting.

1 Introduction

Relation Extraction (RE) refers to classifying semantic relationships between entities within text into predefined types. Conventional RE tasks assume all relations are present at once, ignoring the fact that new relations continually emerge in the real world. Few-shot Continual Relation Extraction (FCRE) is a subfield of continual learning (Hai et al. 2024; Van et al. 2022; Phan et al. 2022; Tran et al. 2024a,b; Le et al. 2024a) where a model must continually assimilate new emerging relations while avoiding the forgetting of old ones, a task made even more challenging by the limited training data available. The importance of FCRE stems from its relevance to dynamic real-world applications, garnering increas-

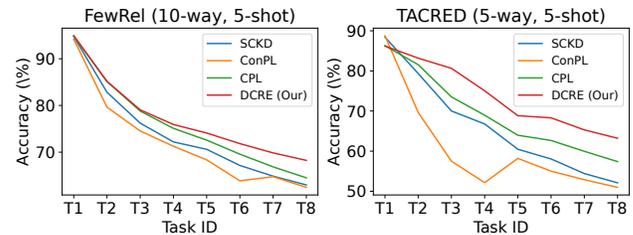


Figure 1: Existing FCRE methods face catastrophic forgetting due to the limited and poor quality of old training samples stored in the memory buffer.

ing interest in the field (Chen, Wu, and Shi 2023a; Le et al. 2024c, 2025).

State-of-the-art approaches to FCRE often rely on memory-based methods for continual learning (Lopez-Paz and Ranzato 2017; Nguyen et al. 2023; Le et al. 2024b; Dao et al. 2024). However, these methods frequently suffer from overfitting to the limited samples stored in memory buffers. This overfitting hampers the reinforcement of previously learned knowledge, leading to catastrophic forgetting—a marked decline in performance on learnt relations when new ones are introduced (Figure 1). The few-shot scenario of FCRE exacerbates these issues, as the scarcity of data not only impedes learning on new tasks, but also hinders helpful data augmentation, which are crucial in many methods (Shin et al. 2017).

In order to improve on these methods, we must not completely disregard them or dwell on their weaknesses, but rather contemplate their biggest strength. *Why do so many methods use the memory buffer in the first place?* The primary objective of these replay buffers is to rehearse and reinforce past knowledge, providing the model with something to "look back" at during training. However, these past samples may not always be representative of the entire class and can still lead to sub-optimal performance. Based on this observation, we propose a straightforward: besides relying

*These authors contributed equally.

†Corresponding Author

on potentially unrepresentative past samples, we leverage our knowledge of the past relations themselves. This insight leads to our approach of generating detailed descriptions for each relation. These descriptions inherently represent the class more accurately than the underlying information from a set of samples, serving as stable pivots for the model to align with past knowledge while learning new information. By using these descriptions, we create a more robust and effective method for Few-Shot Continual Relation Extraction, ensuring better retention of knowledge across tasks.

Overall, our paper makes the following contributions:

- a. We introduce an innovative approach to Few-Shot Continual Relation Extraction that leverages Large Language Models (LLMs) to generate comprehensive descriptions for each relation. These descriptions serve as stable class representations in the latent space during training. Unlike the variability and limitations of a limited set of samples from the memory buffer, these descriptions define the inherent meaning of the relations, offer a more reliable anchor, significantly reducing the risk of catastrophic forgetting. Importantly, LLMs are employed exclusively for generating descriptions and do not participate in the training or inference processes, ensuring that our method incurs minimal computational overhead.
- b. We design a bi-encoder retrieval learning framework for both sample and class representation learning. In addition to sample representation contrastive learning, we integrate a description-pivot learning process, ensuring alignment of samples which maximize their respective class samples proximity, while non-matching samples are distanced.
- c. Building on the enhanced representations, we introduce the *Descriptive Retrieval Inference* (DRI) strategy. In this approach, each sample "retrieves" the most fitting relation using a reciprocal rank fusion score that integrates both class descriptions and class prototypes, effectively finalizing the retrieval-based paradigm that underpins our method.

2 Background

2.1 Problem Formulation

In Few-Shot Continual Relation Extraction (FCRE), a model must continuously assimilate new knowledge from a sequential series of tasks. For each t -th task, the model undergoes training on the dataset $D^t = \{(x_i^t, y_i^t)\}_{i=1}^{N \times K}$. Here, N represents the number of relations in the task R^t , and K denotes the limited number of samples per relation, reflecting the few-shot learning scenario. Each sample (x, y) includes a sentence x containing a pair of entities (e_h, e_t) and a relation label $y \in R$. This type of task setup is referred to as " N -way- K -shot" (Chen, Wu, and Shi 2023a). Upon completion of task t , the dataset D^t should not be extensively included in subsequent learning, as continual learning aims to avoid retraining on all prior data. Ultimately, the model's performance is assessed on a test set which encompasses all encountered relations $\tilde{R}^T = \bigcup_{t=1}^T R^t$.

For clarity, each task in FCRE can be viewed as a conventional relation extraction problem, with the key challenge being the scarcity of samples available for learning. The primary goal of FCRE is to develop a model that can consistently acquire new knowledge from limited data while retaining competence in previously learned tasks. In the following subsections, we will explore the key aspects of FCRE models as addressed by state-of-the-art studies.

2.2 Encoding Latent Representation

A key initial consideration in Relation Extraction is how to formalize the latent representation of the input, as the output of a Transformer (Vaswani et al. 2017) is a matrix. In this work, we adopt a method recently introduced by Ma et al. (2024). Given an input sentence x , which includes a head entity e_h and a tail entity e_t , we reformulate it into a Cloze-style phrase $T(x)$ by incorporating a [MASK] token, which represents the relation between the entities. Specifically, the template is structured as follows:

$$T(x) = x [v_{0:n_0-1}] e_h [v_{n_0:n_1-1}] [\text{MASK}] [v_{n_1:n_2-1}] e_t [v_{n_2:n_3-1}]. \quad (1)$$

Each $[v_i]$ denotes a learnable continuous token, and n_j determines the number of tokens in each phrase. In our specific implementation, we use BERT's [UNUSED] tokens as $[v]$. The soft prompt phrase length is set to 3 tokens, meaning n_0, n_1, n_2 and n_3 correspond to the values of 3, 6, 9, and 12, respectively. We then forward the templated sentence $T(x)$ through BERT to encode it into a sequence of continuous vectors, from which we obtain the hidden representation z of the input, corresponding to the position of the [MASK] token:

$$z = [\mathcal{M} \circ T](x)[\text{position}([\text{MASK}])], \quad (2)$$

where \mathcal{M} denotes the backbone pre-trained language model. This latent representation is then passed through an MLP for prediction, enabling the model to learn which relation that best fills the [MASK] token.

2.3 Learning Latent Representation

In conventional Relation Extraction scenarios, a basic framework typically employs a backbone PLM followed by an MLP classifier to directly map the input space to the label space using Cross Entropy Loss. However, this approach proves inadequate in data-scarce settings (Snell, Swersky, and Zemel 2017). Consequently, training paradigms which directly target the latent space, such as contrastive learning, emerge as more suitable approaches. To enhance the semantic richness of the information extracted from the training samples, two popular losses are often utilized: *Supervised Contrastive Loss* and *Hard Soft Margin Triplet Loss*.

Supervised Contrastive Loss. To enhance the model's discriminative capability, we employ the Supervised Contrastive Loss (SCL) (Khosla et al. 2020). This loss function is designed to bring positive pairs of samples, which share the same class label, closer together in the latent space. Simultaneously, it pushes negative pairs, belonging to different classes, further apart. Let z_x represent the hidden vector

You are a professional data scientist, working in a relation extraction project.

Given a relation and its description, you are asked to write a more detailed description of the relation and provide 3 sentence examples of the relation.

The relation is: `relation_name`

The description is: `raw_description`

Please generate `K` diverse samples of (relation description, examples).

Your response:

Figure 2: Prompt to generate relation descriptions with LLMs.

output of sample x , the positive pairs (z_x, z_p) are those who share a class, while the negative pairs (z_x, z_n) correspond to different labels. The SCL is computed as follows:

$$\mathcal{L}_{\text{SC}}(x) = - \sum_{p \in P(x)} \log \frac{f(z_x, z_p)}{\sum_{u \in \mathcal{D} \setminus \{x\}} f(z_x, z_u)} \quad (3)$$

where $f(\mathbf{x}, \mathbf{y}) := \exp\left(\frac{\gamma(\mathbf{x}, \mathbf{y})}{\tau}\right)$, $\gamma(\cdot, \cdot)$ denotes the cosine similarity function, and τ is the temperature scaling hyperparameter. $P(x)$ and \mathcal{D} denote the sets of positive samples with respect to sample x and the training set, respectively.

Hard Soft Margin Triplet Loss. To achieve a balance between flexibility and discrimination, the Hard Soft Margin Triplet Loss (HSMT) integrates both hard and soft margin triplet loss concepts (Hermans, Beyer, and Leibe 2017). This loss function is designed to maximize the separation between the most challenging positive and negative samples, while preserving a soft margin for improved flexibility. Formally, the loss is defined as:

$$\mathcal{L}_{\text{ST}}(x) = - \log \left(1 + \max_{p \in P(x)} e^{\xi(z_x, z_p)} - \min_{n \in N(x)} e^{\xi(z_x, z_n)} \right), \quad (4)$$

where $\xi(\cdot, \cdot)$ denotes the Euclidean distance function. The objective of this loss is to ensure that the hardest positive sample is as distant as possible from the hardest negative sample, thereby enforcing a flexible yet effective margin.

During training, these two losses is aggregated and referred to as the *Sample-based learning loss*:

$$\mathcal{L}_{\text{Samp}} = \beta_{\text{SC}} \cdot \mathcal{L}_{\text{SC}} + \beta_{\text{ST}} \cdot \mathcal{L}_{\text{ST}} \quad (5)$$

3 Proposed Method

3.1 Label Descriptions

A core component of our method is achieving robust class latent representations, making class encoding crucial. To this end, having detailed definitions for each label, alongside the hidden information extracted from the samples, is essential for our approach. In fact, the datasets used for benchmarking already provide each relation with a concise description,

which we refer to as the *Raw description*. While leveraging these descriptions has shown promise in previous work (Luo et al. 2024), this approach remains limited due to its reliance on a one-to-one mapping between input embeddings and a single label description representation per task. This singular approach fails to offer rich, diverse, and robust information about the labels, leading to potential noise, instability, and suboptimal model performance.

To address these limitations, we employ Gemini 1.5 (Team et al. 2023; Reid et al. 2024) to generate K diverse, detailed, and illustrative descriptions for each relation. In particular, for each label, the respective raw description will be fed into the *LLM prompt*, serving as an expert-in-the-loop to guide the model. Our prompt template is depicted in Figure 2.

3.2 Description-pivot Learning

The single most valuable quality of class descriptions in our problem is that they are literal definitions of a relation, which makes them more accurate representations of that class than the underlying information from a set of samples. Thanks to this strength, they serve as stable knowledge anchors for the model to rehearse from, enabling effective reinforcement of old knowledge while assimilating new information. Unlike the variability of individual samples, a description remains consistent, providing a more reliable reference point for the model to rehearse from, effectively mitigating catastrophic forgetting.

To fully leverage this inherent advantage, we integrate these descriptions into the training process, framing the task as one of retrieving definition, which embodies real-world meaning, rather than a straightforward categorical classification. By doing so, we capitalize on the unchanging nature of descriptions, making them the focal point of our model’s learning. Specifically, we incorporate two description-centric losses to enhance this retrieval-oriented approach:

$$\mathcal{L}_{\text{Des}} = \beta_{\text{HM}} \cdot \mathcal{L}_{\text{HM}} + \beta_{\text{MI}} \cdot \mathcal{L}_{\text{MI}}. \quad (6)$$

Here, \mathcal{L}_{HM} and \mathcal{L}_{MI} denote the Hard Margin Loss and the Mutual Information Loss, respectively. These losses are elaborated upon in the following paragraphs.

Hard Margin Loss. The Hard Margin Loss leverages label descriptions to refine the model’s ability to distinguish between hard positive and hard negative pairs. Given the output hidden vectors $\{d_x^k\}_{k=1, \dots, K}$ from BERT corresponding to the label description of sample x , and z_p and z_n representing the hidden vectors of positive and negative samples respectively, the loss function is formulated to maximize the alignment between d_x^k and its corresponding positive sample, while enforcing a strict margin against negative samples.

Specifically, the loss is formulated as follows:

$$\mathcal{L}_{\text{HM}}(x) = \sum_{k=1}^K \mathcal{L}_{\text{HM}}^k(x), \quad (7)$$

$$\begin{aligned} \mathcal{L}_{\text{HM}}^k(x) = & \sum_{p \in P_{\text{H}}(x)} (1 - \gamma(\mathbf{d}_x^k, \mathbf{z}_p))^2 \\ & + \sum_{n \in N_{\text{H}}(x)} \max(0, m - 1 + \gamma(\mathbf{d}_x^k, \mathbf{z}_n))^2, \end{aligned} \quad (8)$$

where m is a margin hyperparameter; $\gamma(\cdot, \cdot)$ denotes the cosine similarity function; $P_{\text{H}}(x)$ and $N_{\text{H}}(x)$ represent the sets of hard positive and hard negative samples, respectively. They are determined by comparing the similarity between \mathbf{d}_x^k and both positive and negative pairs, specifically focusing on the most challenging pairs where the similarity to negative samples is close to or greater than that of positive samples, defined as follows:

$$\begin{aligned} P_{\text{H}}(x) = & \{p \in P(x) | 1 - \gamma(\mathbf{d}_x^k, \mathbf{z}_p) \\ & > \min_{n \in N(x)} (1 - \gamma(\mathbf{d}_x^k, \mathbf{z}_n)), \forall k \in [K]\}, \end{aligned} \quad (9)$$

$$\begin{aligned} N_{\text{H}}(x) = & \{n \in N(x) | 1 - \gamma(\mathbf{d}_x^k, \mathbf{z}_n) \\ & < \max_{p \in P(x)} (1 - \gamma(\mathbf{d}_x^k, \mathbf{z}_p)), \forall k \in [K]\}. \end{aligned} \quad (10)$$

By utilizing the label description vectors $\{\mathbf{d}_x^k\}$, optimizing $\mathcal{L}_{\text{HM}}(x)$ effectively sharpens the model’s decision boundary, reducing the risk of confusion between similar classes and improving overall performance in few-shot learning scenarios. The loss penalizes the model more heavily for misclassifications involving these hard samples, ensuring that the model pays particular attention to the most difficult cases, thereby enhancing its discriminative power.

Mutual Information Loss. The Mutual Information (MI) Loss is designed to maximize the mutual information between the input sample’s hidden representation \mathbf{z}_x of x and its corresponding retrieved descriptions, promoting a more informative alignment between them. Let \mathbf{d}_n be a hidden vector of other label descriptions than x . According to van den Oord, Li, and Vinyals (2018), the Mutual Information $MI(x)$ between the input embedding \mathbf{z}_x and its corresponding label description follows the following inequation:

$$MI \geq \log B + \text{InfoNCE}(\{x_i\}_{i=1}^B; h), \quad (11)$$

where we have defined:

$$\begin{aligned} \text{InfoNCE}(\{x_i\}_{i=1}^B; h) = & \frac{1}{B} \sum_{i=1}^B \log \frac{\sum_{k=1}^K h(\mathbf{z}_i, \mathbf{d}_i^k)}{\sum_{j=1}^B \sum_{k=1}^K h(\mathbf{z}_j, \mathbf{d}_j^k)}, \end{aligned} \quad (12)$$

where $h(\mathbf{z}_j, \mathbf{d}_j^k) = \exp\left(\frac{\mathbf{z}_j^T W \mathbf{d}_j^k}{\tau}\right)$. Here, τ is the temperature, B is mini-batch size and W is a trainable parameter.

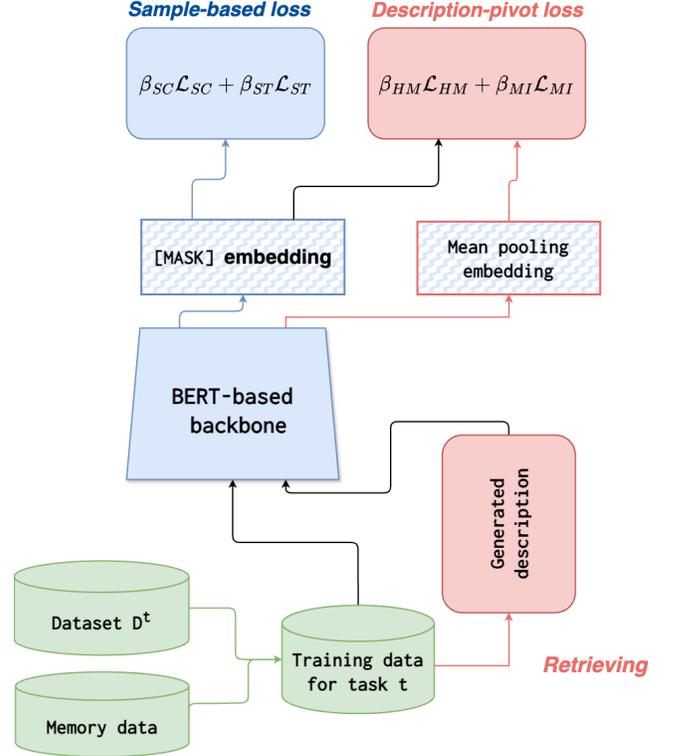


Figure 3: Our Framework.

Finally, the MI loss function in our implementation is:

$$\begin{aligned} \mathcal{L}_{\text{MI}}(x) = & -\log \frac{\sum_{k=1}^K h(\mathbf{z}_x, \mathbf{d}_x^k)}{\sum_{k=1}^K h(\mathbf{z}_x, \mathbf{d}_x^k) + \sum_{n \in N(x)} \sum_{k=1}^K h(\mathbf{z}_x, \mathbf{d}_n^k)} \end{aligned} \quad (13)$$

This loss ensures that the representation of the input sample is strongly associated with its corresponding label, while reducing its association with incorrect labels, thereby enhancing the discriminative power of the model.

Joint Training Objective Function. Our model is trained using a combination of the *Sample-based learning loss* mentioned in Section 2.3 and our description-pivot loss \mathcal{L}_{Des} , weighted by their respective coefficients:

$$\begin{aligned} \mathcal{L}(x) = & \mathcal{L}_{\text{Samp}} + \mathcal{L}_{\text{Des}} \\ = & \beta_{\text{SC}} \cdot \mathcal{L}_{\text{SC}}(x) + \beta_{\text{ST}} \cdot \mathcal{L}_{\text{ST}}(x) \\ & + \beta_{\text{HM}} \cdot \mathcal{L}_{\text{HM}}(x) + \beta_{\text{MI}} \cdot \mathcal{L}_{\text{MI}}(x), \end{aligned} \quad (14)$$

where β_{SC} , β_{ST} , β_{HM} , and β_{MI} are hyperparameters. This joint objective enables the model to leverage the strengths of each individual loss, facilitating robust and effective learning in Few-Shot Continual Relation Extraction tasks.

Training Procedure. Algorithm 1 outlines the end-to-end training process at each task \mathcal{T}^j , with Φ_{j-1} denoting the model after training on the previous $j - 1$ tasks. In line with

FewRel (10-way-5-shot)

Method	\mathcal{T}^1	\mathcal{T}^2	\mathcal{T}^3	\mathcal{T}^4	\mathcal{T}^5	\mathcal{T}^6	\mathcal{T}^7	\mathcal{T}^8	$\Delta \downarrow$
RP-CRE	93.97 \pm 0.64	76.05 \pm 2.36	71.36 \pm 2.83	69.32 \pm 3.98	64.95 \pm 3.09	61.99 \pm 2.09	60.59 \pm 1.87	59.57 \pm 1.13	34.40
CRL	94.68 \pm 0.33	80.73 \pm 2.91	73.82 \pm 2.77	70.26 \pm 3.18	66.62 \pm 2.74	63.28 \pm 2.49	60.96 \pm 2.63	59.27 \pm 1.32	35.41
CRECL	93.93 \pm 0.22	82.55 \pm 6.95	74.13 \pm 3.59	69.33 \pm 3.87	66.51 \pm 4.05	64.60 \pm 1.92	62.97 \pm 1.46	59.99 \pm 0.65	33.94
ERDA	92.43 \pm 0.32	64.52 \pm 2.11	50.31 \pm 3.32	44.92 \pm 3.77	39.75 \pm 3.34	36.36 \pm 3.12	34.34 \pm 1.83	31.96 \pm 1.91	60.47
SCKD	94.77 \pm 0.35	82.83 \pm 2.61	76.21 \pm 1.61	72.19 \pm 1.33	70.61 \pm 2.24	67.15 \pm 1.96	64.86 \pm 1.35	62.98 \pm 0.88	31.79
ConPL**	95.18 \pm 0.73	79.63 \pm 1.27	74.54 \pm 1.13	71.27 \pm 0.85	68.35 \pm 0.86	63.86 \pm 2.03	64.74 \pm 1.39	62.46 \pm 1.54	32.72
CPL	94.87	85.14	78.80	75.10	72.57	69.57	66.85	64.50	30.37
CPL + MI	94.69 \pm 0.7	85.58 \pm 1.88	80.12 \pm 2.45	75.71 \pm 2.28	73.90 \pm 1.8	70.72 \pm 0.91	68.42 \pm 1.77	66.27 \pm 1.58	28.42
DCRE	94.93 \pm 0.39	85.14 \pm 2.27	79.06 \pm 1.68	75.92 \pm 2.03	74.10 \pm 2.53	71.83 \pm 2.17	69.84 \pm 1.48	68.24 \pm 0.79	26.69

TACRED (5-way-5-shot)

Method	\mathcal{T}^1	\mathcal{T}^2	\mathcal{T}^3	\mathcal{T}^4	\mathcal{T}^5	\mathcal{T}^6	\mathcal{T}^7	\mathcal{T}^8	$\Delta \downarrow$
RP-CRE	87.32 \pm 1.76	74.90 \pm 6.13	67.88 \pm 4.31	60.02 \pm 5.37	53.26 \pm 4.67	50.72 \pm 7.62	46.21 \pm 5.29	44.48 \pm 3.74	42.84
CRL	88.32 \pm 1.26	76.30 \pm 7.48	69.76 \pm 5.89	61.93 \pm 2.55	54.68 \pm 3.12	50.92 \pm 4.45	47.00 \pm 3.78	44.27 \pm 2.51	44.05
CRECL	87.09 \pm 2.50	78.09 \pm 5.74	61.93 \pm 4.89	55.60 \pm 5.78	53.42 \pm 2.99	51.91 \pm 2.95	47.55 \pm 3.38	45.53 \pm 1.96	41.56
ERDA	81.88 \pm 1.97	53.68 \pm 6.31	40.36 \pm 3.35	36.17 \pm 3.65	30.14 \pm 3.96	22.61 \pm 3.13	22.29 \pm 1.32	19.42 \pm 2.31	62.46
SCKD	88.42 \pm 0.83	79.35 \pm 4.13	70.61 \pm 3.16	66.78 \pm 4.29	60.47 \pm 3.05	58.05 \pm 3.84	54.41 \pm 3.47	52.11 \pm 3.15	36.31
ConPL**	88.77 \pm 0.84	69.64 \pm 1.93	57.50 \pm 2.48	52.15 \pm 1.59	58.19 \pm 2.31	55.01 \pm 3.12	52.88 \pm 3.66	50.97 \pm 3.41	37.80
CPL	86.27	<u>81.55</u>	73.52	<u>68.96</u>	63.96	62.66	59.96	57.39	28.88
CPL + MI	85.67 \pm 0.8	82.54 \pm 2.98	75.12 \pm 3.67	70.65 \pm 2.75	66.79 \pm 2.18	65.17 \pm 2.48	61.25 \pm 1.52	59.48 \pm 3.53	26.19
DCRE	86.20 \pm 1.35	83.18 \pm 8.04	80.65 \pm 3.06	75.05 \pm 3.07	68.83 \pm 5.05	68.30 \pm 4.28	65.30 \pm 2.74	63.21 \pm 2.39	22.99

Table 1: Accuracy (%) of methods using BERT-based backbone after training for each task. The best results are in **bold**. **Results of ConPL are reproduced

memory-based continual learning methods, we maintain a memory buffer \tilde{M}_{j-1} that stores a few representative samples from all previous tasks $\mathcal{T}^1, \dots, \mathcal{T}^{j-1}$, along with a relation description set \tilde{E}_{j-1} that holds the descriptions of all previously encountered relations.

- Initialization** (Line 1–2): The model for the current task, Φ_j , is initialized with the parameters of Φ_{j-1} . We update the relation description set \tilde{E}_j by incorporating new relation descriptions from E_j .
- Training on the Current Task** (Line 3): We train Φ_j on D_j to learn the novel relations introduced in \mathcal{T}^j .
- Memory Update** (Lines 4–8): We select L representative samples from D_j for each relation $r \in R_j$. These are the L samples whose latent representations are closest to the 1-means centroid of all class samples. These samples constitute the memory M_r , leading to an updated overall memory $\tilde{M}_j = \tilde{M}_{j-1} \cup M_j$ and an updated relation set $\tilde{R}_j = \tilde{R}_{j-1} \cup R_j$.
- Prototype Storing** (Line 9): A prototype set \tilde{P}_j is generated based on the updated memory \tilde{M}_j . We generate a prototype set \tilde{P}_j based on the updated memory \tilde{M}_j .
- Memory Training** (Line 10): We refine Φ_j by training on the augmented memory set \tilde{M}_j^* , ensuring that the model preserves knowledge of relations from previous tasks.

Algorithm 1: Training procedure at each task \mathcal{T}^j

Input: $\Phi_{j-1}, \tilde{R}_{j-1}, \tilde{M}_{j-1}, \tilde{K}_{j-1}, D_j, R_j, K_j$.

Output: $\Phi_j, \tilde{M}_j, \tilde{K}_j, \tilde{P}_j$.

- Initialize Φ_j from Φ_{j-1}
- $\tilde{K}_j \leftarrow \tilde{K}_{j-1} \cup K_j$
- Update Φ_j by L on D_j (train on current task)
- $\tilde{M}_j \leftarrow \tilde{M}_{j-1}$
- for** each $r \in R_j$ **do**
- pick L samples in D_j and add them into \tilde{M}_j
- end for**
- $\tilde{R}_j \leftarrow \tilde{R}_{j-1} \cup R_j$
- Update \tilde{P}_j with new data in D_j (for inference)
- Update Φ_j by \mathcal{L} on \tilde{M}_j and D_j^* (train on memory)

3.3 Descriptive Retrieval Inference

Traditional methods such as Nearest Class Mean (NCM) (Ma et al. 2024) predict relations by selecting the class whose prototype has the smallest distance to the test sample x . While effective, this approach relies solely on distance metrics, which may not fully capture the nuanced relationships between a sample and the broader context provided by class descriptions.

Rather than merely seeking the closest prototype, we aim to retrieve the class description that best aligns with the input, thereby leveraging the inherent semantic meaning of the label. To achieve this, we introduce *Descriptive Retrieval In-*

FewRel (10-way-5-shot)

Method	\mathcal{T}^1	\mathcal{T}^2	\mathcal{T}^3	\mathcal{T}^4	\mathcal{T}^5	\mathcal{T}^6	\mathcal{T}^7	\mathcal{T}^8	$\Delta \downarrow$
CPL	97.25 \pm 0.30	89.29 \pm 2.51	85.56 \pm 1.21	82.10 \pm 2.02	79.96 \pm 2.72	78.41 \pm 3.22	76.42 \pm 2.25	75.20 \pm 2.33	22.05
DCRE	96.92 \pm 0.16	88.95 \pm 1.72	87.12 \pm 1.52	85.44 \pm 1.91	84.89 \pm 2.12	83.52 \pm 1.46	81.64 \pm 0.69	80.34 \pm 0.55	16.58

TACRED (5-way-5-shot)

Method	\mathcal{T}^1	\mathcal{T}^2	\mathcal{T}^3	\mathcal{T}^4	\mathcal{T}^5	\mathcal{T}^6	\mathcal{T}^7	\mathcal{T}^8	$\Delta \downarrow$
CPL	88.74 \pm 0.44	85.16 \pm 5.38	78.35 \pm 4.46	77.50 \pm 4.04	76.01 \pm 5.04	76.30 \pm 4.41	74.51 \pm 5.06	73.83 \pm 4.91	14.91
DCRE	89.06 \pm 0.59	87.41 \pm 5.54	84.91 \pm 3.38	84.18 \pm 2.44	82.74 \pm 3.64	81.92 \pm 2.33	79.34 \pm 2.89	79.10 \pm 2.37	9.96

Table 2: Accuracy (%) of methods using LLM2Vec-based backbone after training for each task. The best results are in **bold**.

ference (DRI), a retrieval mechanism fusing two distinct reciprocal ranking scores. This approach not only considers the proximity of a sample to class prototypes but also incorporates cosine similarity measures between the sample’s hidden representation \mathbf{z} and relation descriptions generated by an LLM. This dual focus on both spatial and semantic alignment ensures that the final prediction is informed by a richer, more robust understanding of the relations.

Given a sample x with hidden representation \mathbf{z} and a set of relation prototypes $\{\mathbf{p}_r\}_{r=1}^n$, the inference process begins by calculating the negative Euclidean distance between \mathbf{z} and each prototype \mathbf{p}_r :

$$\mathbf{E}(x, r) = -\|\mathbf{z} - \mathbf{p}_r\|_2, \quad (16)$$

$$\mathbf{p}_r = \frac{1}{L} \sum_{i=1}^L \mathbf{z}_i, \quad (17)$$

where L is the memory size per relation. Simultaneously, we compute the cosine similarity between the hidden representation and each relation description prototype, $\gamma(\mathbf{z}, \mathbf{d}_r)$. These two scores are combined into DRI score of sample x w.r.t relation r for inference, ensuring that predictions align with both label prototypes and relation descriptions:

$$\text{DRI}(x, r) = \frac{\alpha}{\epsilon + \text{rank}(\mathbf{E}(x, r))} + \frac{1 - \alpha}{\epsilon + \text{rank}(\gamma(\mathbf{z}, \mathbf{d}_r))}, \quad (18)$$

where $d_r = \frac{1}{K} \sum_{i=1}^K d_r^i$, $\text{rank}(\cdot)$ represents the rank position of the score among all relations. The α hyperparameter balances the contributions of the Euclidean distance-based score and the cosine similarity score in the final ranking for inference, and ϵ is a hyperparameter that controls the influence of lower-ranked relations in the final prediction. By adjusting ϵ , we can fine-tune the model’s sensitivity to less prominent relations. Finally, the predicted relation label y^* is predicted as the one corresponding to the highest DRI score:

$$y_x^* = \underset{r=1, \dots, n}{\text{argmax}} \text{DRI}(x, r) \quad (19)$$

This fusion approach for inference complements the learning paradigm, ensuring consistency and reliability throughout the FCRE process. By effectively balancing the strengths of prototype-based proximity and description-based semantic similarity, it leads to more accurate and robust predictions across sequential tasks.

4 Experiments

4.1 Settings

We conduct experiments using two pre-trained language models, BERT (Devlin et al. 2019) and LLM2Vec (BehnamGhader et al. 2024), on two widely used benchmark datasets for Relation Extraction: FewRel (Han et al. 2018) and TACRED (Zhang et al. 2017). We benchmark our methods against state-of-the-art baselines: **SCKD** (Wang, Wang, and Hu 2023), **RP-CRE** (Cui et al. 2021), **CRL** (Zhao et al. 2022), **CRECL** (Hu et al. 2022), **ERDA** (Qin and Joty 2022), **ConPL** (Chen, Wu, and Shi 2023b), **CPL** (Ma et al. 2024), **CPL+MI** (Tran et al. 2024c).

4.2 Experiment results

Our proposed method yields state-of-the-art accuracy. Table 1 presents the results of our method and the baselines, all using the same pre-trained BERT-based backbone. Our method consistently outperforms all baselines across the board. The performance gap between our method and the strongest baseline, CPL, reaches up to 3.74% on FewRel and 5.82% on TACRED.

To further validate our model, we tested it on LLM2Vec, which provides stronger representation learning than BERT. As shown in Table 2, our model again surpasses CPL, with accuracy drops of only 16.58% on FewRel and 9.96% on TACRED.

These results highlight the effectiveness of our method in leveraging semantic information from descriptions, which helps mitigate forgetting and overfitting, ultimately leading to significant performance improvements.

Exploiting additional descriptions significantly enhances representation learning. Figure 4 presents t-SNE visualizations of the latent space of relations without (left) and with (right) the use of descriptions during training. The visualizations reveal that incorporating descriptions markedly improves the quality of the model’s representation learning. For instance, the brown-orange and purple-green class pairs, which are closely clustered and prone to misclassification in the left image, are more distinctly separated in the right image. Additionally, Figure 5 illustrates that our strategy, which leverages refined descriptions, captures more semantic knowledge related to the labels than the approach using raw descriptions. This advantage bridges the gap imposed by the challenges of few-shot continual learning scenarios,

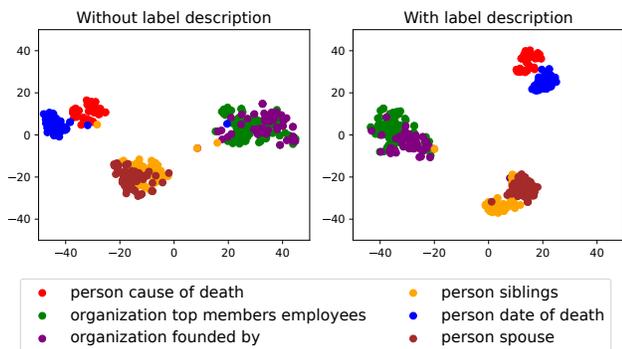


Figure 4: t-SNE visualization of the representations of 6 relations post-training, with and without descriptions, using our retrieval strategy.

leading to superior performance. Figure 6 shows the performance of our model on TACRED as the number of generated expert descriptions per training varies. The results indicate that the model performance generally improves from $K = 3$ and peaks at $K = 7$.

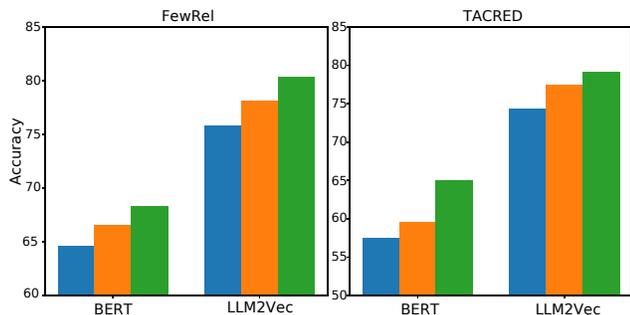


Figure 5: The impact of refined descriptions generated by LLMs. The green, orange, and blue bars show respectively the final accuracies of DCRE when using refined descriptions, original descriptions, and without using descriptions.

Our retrieval-based prediction strategy notably enhances model performance. Table 3 demonstrates that by leveraging the rich information from generated descriptions, our proposed strategy improves the model’s performance by up to 1.31% on FewRel and 6.66% on TACRED compared to traditional NCM-based classification. The harmonious integration of NCM-based prototype proximity and description-based semantic similarity enables our strategy to deliver more accurate and robust predictions across sequential tasks.

4.3 Ablation study

Table 4 present evaluation results that closely examine the role of each component in the objective function during training. The findings underscore the critical importance of \mathcal{L}_{MI} and \mathcal{L}_{HM} , both of which leverage instructive descriptions from LLMs, aided by *Raw descriptions*. Because when

Method	FewRel		TACRED	
	BERT	LLM2Vec	BERT	LLM2Vec
NCM	66.93	79.26	58.26	75.00
DRI (Ours)	68.24	80.34	63.21	79.10

Table 3: DRI and NCM prediction.

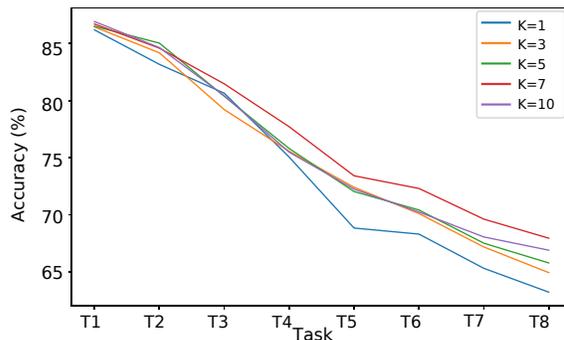


Figure 6: Model performance when varying K , on TACRED 5-way 5-shot.

we ablate one of them, the final accuracy can be reduced by 6% on the BERT-based model, and 10% on the LLM2VEC-based model.

Method	BERT		LLM2Vec	
	FewRel	TACRED	FewRel	TACRED
DCRE (Our)	68.24	63.21	80.34	79.10
w/o \mathcal{L}_{SC}	67.58	62.11	78.39	77.01
w/o \mathcal{L}_{MI}	65.10	57.23	70.61	74.17
w/o \mathcal{L}_{HM}	66.20	62.46	77.22	74.75
w/o \mathcal{L}_{ST}	67.54	59.56	77.48	73.77

Table 4: Ablation study.

5 Conclusion

In this work, we propose a novel retrieval-based approach to address the challenging problem of Few-shot Continual Relation Extraction. By leveraging large language models to generate rich relation descriptions, our bi-encoder training paradigm enhances both sample and class representations and also enables a robust retrieval-based prediction method that maintains performance across sequential tasks. Extensive experiments demonstrate the effectiveness of our approach in advancing the state-of-the-art and overcoming the limitations of traditional memory-based techniques, underscoring the potential of language models and retrieval techniques for dynamic real-world relationship identification.

References

- BehnamGhader, P.; Adlakha, V.; Mosbach, M.; Bahdanau, D.; Chapados, N.; and Reddy, S. 2024. LLM2Vec: Large Language Models Are Secretly Powerful Text Encoders. *arXiv preprint arXiv:2404.05961*.
- Chen, X.; Wu, H.; and Shi, X. 2023a. Consistent Prototype Learning for Few-Shot Continual Relation Extraction. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7409–7422. Toronto, Canada: Association for Computational Linguistics.
- Chen, X.; Wu, H.; and Shi, X. 2023b. Consistent Prototype Learning for Few-Shot Continual Relation Extraction. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, 7409–7422. Association for Computational Linguistics.
- Cui, L.; Yang, D.; Yu, J.; Hu, C.; Cheng, J.; Yi, J.; and Xiao, Y. 2021. Refining Sample Embeddings with Relation Prototypes to Enhance Continual Relation Extraction. 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)*, 232–243. Online: Association for Computational Linguistics.
- Dao, V.; Pham, V.-C.; Tran, Q.; Le, T.-T.; Ngo, L.; and Nguyen, T. 2024. Lifelong Event Detection via Optimal Transport. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 12610–12621.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Hai, N. L.; Nguyen, T.; Van, L. N.; Nguyen, T. H.; and Than, K. 2024. Continual variational dropout: a view of auxiliary local variables in continual learning. *Machine Learning*, 113(1): 281–323.
- Han, X.; Zhu, H.; Yu, P.; Wang, Z.; Yao, Y.; Liu, Z.; and Sun, M. 2018. FewRel: A Large-Scale Supervised Few-Shot Relation Classification Dataset with State-of-the-Art Evaluation. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4803–4809. Brussels, Belgium: Association for Computational Linguistics.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In Defense of the Triplet Loss for Person Re-Identification. *arXiv:1703.07737*.
- Hu, C.; Yang, D.; Jin, H.; Chen, Z.; and Xiao, Y. 2022. Improving Continual Relation Extraction through Prototypical Contrastive Learning. In Calzolari, N.; Huang, C.-R.; Kim, H.; Pustejovsky, J.; Wanner, L.; Choi, K.-S.; Ryu, P.-M.; Chen, H.-H.; Donatelli, L.; Ji, H.; Kurohashi, S.; Paggio, P.; Xue, N.; Kim, S.; Hahm, Y.; He, Z.; Lee, T. K.; Santus, E.; Bond, F.; and Na, S.-H., eds., *Proceedings of the 29th International Conference on Computational Linguistics*, 1885–1895. Gyeongju, Republic of Korea: International Committee on Computational Linguistics.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 18661–18673. Curran Associates, Inc.
- Le, M.; Luu, T. N.; The, A. N.; Le, T.-T.; Nguyen, T.; Nguyen, T. T.; Van, L. N.; and Nguyen, T. H. 2025. Adaptive Prompting for Continual Relation Extraction: A Within-Task Variance Perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Le, M.; Nguyen, A.; Nguyen, H.; Nguyen, T.; Pham, T.; Ngo, L. V.; and Ho, N. 2024a. Mixture of Experts Meets Prompt-Based Continual Learning. In *Advances in Neural Information Processing Systems*.
- Le, T.-T.; Dao, V.; Nguyen, L.; Nguyen, T.-N.; Ngo, L.; and Nguyen, T. 2024b. SharpSeq: Empowering Continual Event Detection through Sharpness-Aware Sequential-task Learning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 3632–3644.
- Le, T.-T.; Nguyen, M.; Nguyen, T. T.; Van, L. N.; and Nguyen, T. H. 2024c. Continual relation extraction via sequential multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18444–18452.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Luo, D.; Gan, Y.; Hou, R.; Lin, R.; Liu, Q.; Cai, Y.; and Gao, W. 2024. Synergistic Anchored Contrastive Pre-training for Few-Shot Relation Extraction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17): 18742–18750.
- Ma, S.; Han, J.; Liang, Y.; and Cheng, B. 2024. Making Pre-trained Language Models Better Continual Few-Shot Relation Extractors. In Calzolari, N.; Kan, M.-Y.; Hoste, V.; Lenci, A.; Sakti, S.; and Xue, N., eds., *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 10970–10983. Torino, Italia: ELRA and ICCL.
- Nguyen, H.; Nguyen, C.; Ngo, L.; Luu, A.; and Nguyen, T. 2023. A spectral viewpoint on continual relation extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 9621–9629.

- Phan, H.; Tuan, A. P.; Nguyen, S.; Linh, N. V.; and Than, K. 2022. Reducing catastrophic forgetting in neural networks via gaussian mixture approximation. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 106–117. Springer.
- Qin, C.; and Joty, S. 2022. Continual Few-shot Relation Learning via Embedding Space Regularization and Data Augmentation. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2776–2789. Dublin, Ireland: Association for Computational Linguistics.
- Reid, M.; Savinov, N.; Telyashin, D.; Lepikhin, D.; Lillcrap, T.; Alayrac, J.-b.; Soricut, R.; Lazaridou, A.; Firat, O.; Schrittwieser, J.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual Learning with Deep Generative Replay. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-shot Learning. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: A Family of Highly Capable Multimodal Models. *arXiv preprint arXiv:2312.11805*.
- Tran, Q.; Phan, H.; Le, M.; Truong, T.; Phung, D.; Ngo, L.; Nguyen, T.; Ho, N.; and Le, T. 2024a. Leveraging Hierarchical Taxonomies in Prompt-based Continual Learning. *arXiv:2410.04327*.
- Tran, Q.; Phan, H.; Tran, L.; Than, K.; Tran, T.; Phung, D.; and Le, T. 2024b. KOPPA: Improving Prompt-based Continual Learning with Key-Query Orthogonal Projection and Prototype-based One-Versus-All. *arXiv:2311.15414*.
- Tran, Q.; Thanh, N.; Anh, N.; Hai, N.; Le, T.; Ngo, L.; and Nguyen, T. 2024c. Preserving Generalization of Language models in Few-shot Continual Relation Extraction. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 13771–13784.
- Van, L. N.; Hai, N. L.; Pham, H.; and Than, K. 2022. Auxiliary local variables for improving regularization/prior approach in continual learning. In *Pacific-Asia conference on knowledge discovery and data mining*, 16–28. Springer.
- van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation Learning with Contrastive Predictive Coding. *CoRR*, abs/1807.03748.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, X.; Wang, Z.; and Hu, W. 2023. Serial Contrastive Knowledge Distillation for Continual Few-shot Relation Extraction. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 12693–12706. Toronto, Canada: Association for Computational Linguistics.
- Zhang, Y.; Zhong, V.; Chen, D.; Angeli, G.; and Manning, C. D. 2017. Position-aware Attention and Supervised Data Improve Slot Filling. In Palmer, M.; Hwa, R.; and Riedel, S., eds., *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 35–45. Copenhagen, Denmark: Association for Computational Linguistics.
- Zhao, K.; Xu, H.; Yang, J.; and Gao, K. 2022. Consistent Representation Learning for Continual Relation Extraction. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Findings of the Association for Computational Linguistics: ACL 2022*, 3402–3411. Dublin, Ireland: Association for Computational Linguistics.