

Consistent Representation Learning for Continual Relation Extraction

Kang Zhao^{1,2}, Hua Xu^{1*}, Jiangong Yang^{1,2}, Kai Gao^{2*}

¹State Key Laboratory of Intelligent Technology and Systems,
Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

²School of Information Science and Engineering,
Hebei University of Science and Technology, Shijiazhuang 050018, China

zhaok7878@gmail.com,
xuhua@tsinghua.edu.cn, yjg219@163.com, gaokai@hebust.edu.cn

Abstract

Continual relation extraction (CRE) aims to continuously train a model on data with new relations while avoiding forgetting old ones. Some previous work has proved that storing a few typical samples of old relations and replaying them when learning new relations can effectively avoid forgetting. However, these memory-based methods tend to overfit the memory samples and perform poorly on imbalanced datasets. To solve these challenges, a consistent representation learning method is proposed, which maintains the stability of the relation embedding by adopting contrastive learning and knowledge distillation when replaying memory. Specifically, supervised contrastive learning based on a memory bank is first used to train each new task so that the model can effectively learn the relation representation. Then, contrastive replay is conducted of the samples in memory and makes the model retain the knowledge of historical relations through memory knowledge distillation to prevent the catastrophic forgetting of the old task. The proposed method can better learn consistent representations to alleviate forgetting effectively. Extensive experiments on FewRel and TACRED datasets show that our method significantly outperforms state-of-the-art baselines and yield strong robustness on the imbalanced dataset. The code is publicly available at <https://github.com/thuiar/CRL>.

1 Introduction

Relation extraction (RE) is an essential issue in information extraction (IE), which can apply to many downstream NLP tasks, such as information retrieval (Xiong et al., 2017) and question and answer (Tao et al., 2018). For example, given a sentence x with the annotated entities pairs e_1 and e_2 , the RE aims to identify the relations between e_1 and e_2 . However, traditional relation extraction models

(Zhou et al., 2016; Soares et al., 2019a) always assume a fixed set of predefined relations and train on a fixed dataset, which cannot handle the growing relation types in real life well.

To solve this situation, continual relation extraction (CRE) is introduced (Wang et al., 2019; Han et al., 2020; Wu et al., 2021; Cui et al., 2021). Compared with traditional relation extraction, CRE aims to help the model learn new relations while maintaining accurate classification of old ones. Wang et al. (2019) shows that continual relation learning needs to alleviate the catastrophic forgetting of old tasks when the model learns new tasks. Because neural networks need to retrain a fixed set of parameters with each training, the most efficient solution to the problem of catastrophic forgetting is to store all the historical data and retrain the model with all the data each time a new relational instance appears. This method can achieve the best effect in continual relation learning, but it is not adopted in real life due to the time and computing power costs.

Some recent works have proposed a variety of methods to alleviate the catastrophic forgetting problem in continual learning, including regularization methods (Kirkpatrick et al., 2017; Zenke et al., 2017; Liu et al., 2018), dynamic architecture methods (Chen et al., 2015; Fernando et al., 2017), and memory-based methods (Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018). Although these methods have been verified in simple image classification tasks, previous works have proved that memory-based methods are the most effective in natural language processing applications (Wang et al., 2019; de Masson D’Autume et al., 2019). In recent years, the memory-based continual relation extraction model has made significant progress in alleviating the problem of catastrophic forgetting (Han et al., 2020; Wu et al., 2021; Cui et al., 2021). Wang et al. (2019) proposes a mechanism for embedding sentence alignment in memory mainte-

*Corresponding Author

nance to ensure the stability of the embedding space. Han et al. (2020) introduces a multi-round joint training process for memory consolidation. But these two methods only explore the problem of catastrophic forgetting in the overall performance of the task sequence. Wu et al. (2021) proposes to integrate curriculum learning. Although it is possible to analyze the characteristics of each subtask and the performance of the corresponding model, it still fails to make full use of the saved sample information. Cui et al. (2021) introduce an attention network to refine the prototype to better recover the interruption of the embedded space. However, this method will produce a bias in the classification of the old task as the new task continues to learn the classifier, which will affect the performance of the old task. Although the above method can alleviate catastrophic forgetting to a certain extent, it does not consider the consistency of relation embedding space.

Because the performance of the model of CRE is sensitive to the quality of sample embedding, it needs to ensure that the learning of new tasks will not damage the embedding of old tasks. Inspired by supervised contrastive Learning (Khosla et al., 2020) to explicitly constrain data embeddings, a consistent representation learning method is proposed for continual relation extraction, which constrains the embedding of old tasks not to occur significantly change through supervised contrastive learning and knowledge distillation. Specifically, the example encoder first trains on the current task data through supervised contrastive learning based on memory bank, and then uses k-means to select representative samples to storage as memory after the training is completed. To relieve catastrophic forgetting, contrastive replay is used to train memorized samples. At the same time, to ensure that the embedding of historical relations does not undergo significant changes, knowledge distillation is used to make the embedding distribution of the new and old tasks consistent. In the testing phase, the nearest class mean (NCM) classifier is used to classify the test sample, which will not be affected by the deviation of the classifier.

In summary, our contributions in this paper are summarized as follows: First, a novel CRE method is proposed, which uses supervised contrastive learning and knowledge distillation to learn consistent relation representations for continual learning. Second, consistent representation learning can en-

sure the stability of the relational embedding space to alleviate catastrophic forgetting and make full use of stored samples. Finally, extensive experiments results on FewRel and TACRED datasets show that the proposed method is better than the latest baseline and effectively mitigates catastrophic forgetting.

2 Related Work

2.1 Continual Learning

Existing continual learning models mainly focus on three areas: (1) Regularization-based methods (Kirkpatrick et al., 2017; Zenke et al., 2017) impose constraints on updating neural weights important to previous tasks for relieving catastrophic forgetting. (2) Dynamic architecture methods (Chen et al., 2015; Fernando et al., 2017) extends the model architecture dynamically to learn new tasks and prevent forgetting old tasks effectively. However, these methods are unsuitable for NLP applications because the model size increases dramatically with increasing tasks. (3) Memory-based methods (Lopez-Paz and Ranzato, 2017; Aljundi et al., 2018; Chaudhry et al., 2018; Mai et al., 2021) saves some samples from old tasks and continuously learns them in new tasks to alleviate catastrophic forgetting. Dong et al. (2021) proposes a simple relational distillation incremental learning framework to balance retaining old knowledge and adapting to new knowledge. Yan et al. (2021) proposes a new two-stage learning method that uses dynamic expandable representation for more effective incremental conceptual modelling. Among these methods, memory-based methods are the most effective in NLP tasks (Wang et al., 2019; Sun et al., 2019; de Masson D’Autume et al., 2019). Inspired by the success of memory-based methods in the field of NLP, we use the framework of memory replay to learn new relations that are constantly emerging.

2.2 Contrastive Learning

Contrastive learning (CL) aims to make the representations of similar samples map closer to each other in the embedded space, while that of dissimilar samples should be farther away (Jaiswal et al., 2021). In recent years, the rise of CL has made great progress in self-supervised representation learning. (Wu et al., 2018; He et al., 2020; Li et al., 2020; Chen and He, 2021). The common point of these works is that no labels are available, so positive and negative pairs were formed through

data augmentations. Recently, supervised contrastive learning (Khosla et al., 2020) has received much attention, which uses label information to extend contrastive learning. Hendrycks and Dietterich (2019) compares the supervised contrastive loss with the cross-entropy loss on the ImageNet-C dataset, and verifies that the supervised contrastive loss is not sensitive to the hyperparameter settings of the optimizer or data enhancement. Chen et al. (2020) proposed a contrastive learning framework for visual representations that does not require a special architecture or memory bank. Khosla et al. (2020) extend the self-supervised batch contrastive approach to the fully-supervised setting, which use supervised contrastive loss learning better representation. Liu and Abbeel (2020) proposed a hybrid discriminant-generative training method based on an energy model. In this paper, contrastive learning is applied to continual relation extraction to extract better relation representation.

3 Methodology

3.1 Problem Formulation

In continual relation extraction, given a series of K tasks $\{T_1, T_2, \dots, T_K\}$, where the k -th task has its own training set D_k and relation set R_k . Each task T_k is a traditional supervised classification task, including a series of examples and their corresponding labels $\{(x_i, y_i)\}_{i=1}^N$, where x_i is the input data, including the natural language text and entity pair, and $y_i \in R_k$ is the relation label. The goal of continual relation learning is to train the model, which keeps learning new tasks while avoiding catastrophic forgetting of previous learning tasks. In other words, after learning the k -th task, the model can identify the relation of a given entity pair into \hat{R}_k , where $\hat{R}_k = \cup_{i=1}^k R_i$ is the relation set already observed till the k -th task.

In order to mitigate catastrophic forgetting in continual relational extraction, episodic memory modules have been used in previous work (Wang et al., 2019; Han et al., 2020; Cui et al., 2021), to store small samples in historical tasks. Inspired by (Cui et al., 2021), we store several representative samples for each relation. Therefore, the episodic memory module for the observed relations in $T_1 \sim T_k$ is $\hat{M}_k = \cup_{r \in \hat{R}_k} M_r$, where $M_r = \{(x_i, y_i)\}_{i=1}^O$, r represents a certain relation, and O is sample number (memory size).

Algorithm 1 Training procedure for T_k

Input:

The training set of D_k of the k -th task, encoder \mathbf{E} , projection head Proj, history memory M_{k-1} , current relation set R_k , history relation set \hat{R}_{k-1}

Output:

encoder $f_k(\cdot)$, history memory M_k , history relation set \hat{R}_k

```

1: if  $T_k$  is not the first task then
2:   get memory knowledge with  $\mathbf{E}$  on  $M_{k-1}$ ;
3: end if
4:  $M_b \leftarrow \mathbf{E}(D_k)$ ;
5: for  $i \leftarrow 1$  to  $epoch_1$  do
6:   for each  $x_j \in D_k$  do
7:     Sample from  $M_b$ ;
8:     Update  $\mathbf{E}$  and Proj with  $\nabla L_{CL}$ ;
9:     Update  $M_b$ ;
10:  end for
11: end for
12: Select informative examples from  $D_k$  to store into  $\hat{M}$ 
13:  $M_k \leftarrow M_{k-1} \cup \hat{M}$ ;
14:  $\hat{R}_k \leftarrow \hat{R}_{k-1} \cup R_k$ ;
15: if  $T_k$  is not the first task then
16:    $\tilde{M}_b \leftarrow \mathbf{E}(M_k)$ ;
17:   for  $i \leftarrow 1$  to  $epoch_2$  do
18:     for each  $x_j \in M_k$  do
19:       Sample from  $\tilde{M}_b$ ;
20:       Update  $\mathbf{E}$  and Proj with  $\nabla L_{CR}$  and  $\nabla L_{KL}$ ;
21:       Update  $\tilde{M}_b$ ;
22:     end for
23:   end for
24:   Select informative examples from  $D_k$  to store into  $\hat{M}$ ;
25:    $M_k \leftarrow M_{k-1} \cup \hat{M}$ 
26: end if
27: return  $\mathbf{E}, M_k, \hat{R}_k$ ;

```

3.2 Framework

The consistent representation learning (CRL) in the current task is described in Algorithm 1, which consists of three main steps: (1) **Init training for new task** (line 4 ~ 11): The parameters of the encoder and projector head are trained on the training sample in D_k with supervised contrastive learning. (2) **Sample selection** (line 12 ~ 13): For each relation $r \in R_k$, we retrieve all samples labeled r from D_k . Then, the k-means algorithm is used

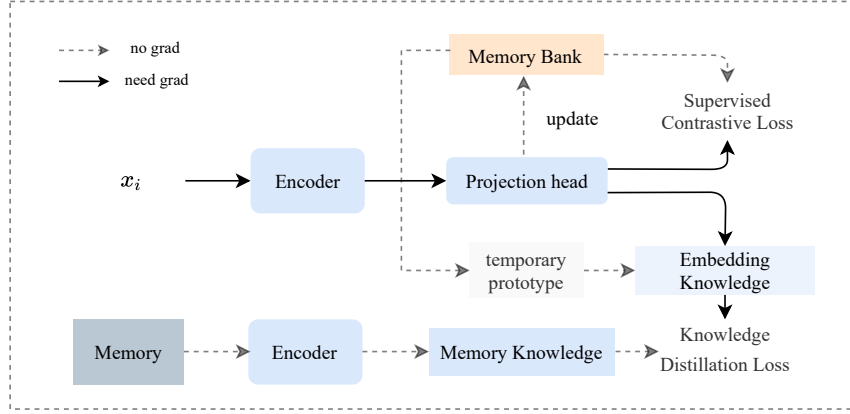


Figure 1: Framework of consistent representation learning.

to cluster the samples. The relation representation of the sample closest to the center is selected and stored in memory for each cluster. (3) **Consistent representation learning** (16 ~ 23): In order to keep the embedding of historical relations in space consistent after learning new tasks, we perform contrastive replay and knowledge distillation constraints on the samples in memory.

3.3 Encoder

The key of CRE is to obtain a better relation representation. The pre-trained language model BERT (Devlin et al., 2019) shows a powerful ability in extracting contextual representation of text. Therefore, BERT is used to encode entity pairs and context information to get the relational representation.

Given a sentence $x = [w_1, \dots, w_{|x|}]$ and a pair of entities (E1, E2), we follow Soares et al. (2019b) augment x with four reserved word pieces to mark the begin and end of each entity mentioned in the sentence. The new token sequence is fed into BERT instead of x . To get the final relation representation between the two entities, the output corresponding to the positions of E1 and E2 are concatenated, and then map it to a high-dimensional hidden representation $\mathbf{h} \in \mathbb{R}^{d_h}$, as follows:

$$\mathbf{h} = \mathbf{W}[\mathbf{h}_{[E1]}; \mathbf{h}_{[E2]}] + \mathbf{b}, \quad (1)$$

where $\mathbf{W} \in \mathbb{R}^{2d_h \times d_h}$ and $\mathbf{b} \in \mathbb{R}^{d_h}$ are trainable parameters. The encoder in which the above-mentioned encoded sentence is a relation representation is denoted as \mathbf{E} .

Then, we use a projection head Proj to obtain the low-dimensional embedding:

$$\tilde{\mathbf{z}} = \text{Proj}(\mathbf{h}), \quad (2)$$

where $\text{Proj}(\cdot) = \text{MLP}(\cdot)$ is composed of two layers of neural networks. The normalized embedding $\mathbf{z} = \tilde{\mathbf{z}} / \|\tilde{\mathbf{z}}\|$ is used for contrastive learning, and the hidden representation is used for classification.

3.4 Initial training for new task

Before training for each new task T_k , we first use Encoder to extract the embedding $\tilde{\mathbf{z}}$ of the relational representation of each sentence in D_k , and use them as the initialized memory bank M_b :

$$M_b \leftarrow \{\mathbf{z}_i\}_{i=1}^N. \quad (3)$$

At the beginning of training, relation representation extraction is performed on each batch B . Then the data embedding is explicitly constrained by clustering through supervised contrastive learning (Khosla et al., 2020):

$$\mathcal{L}_{\text{CL}} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{j \in S_I} \exp(z_i \cdot z_j / \tau)}, \quad (4)$$

where $I = \{1, 2, \dots, |B|\}$ is the set of indices of B . S_I represents the indices set of randomly sampled partial samples from M_b . $P(i) = \{p \in S_I : y_p = y_i\}$ is the indices set that is the same as the z_i label in M_b , and $|P(i)|$ is its cardinality. $\tau \in \mathcal{R}^+$ is an adjustable temperature parameter controlling the separation of classes, the \cdot indicates the dot product.

After backpropagating the gradient of loss on each batch, we update the representation in the memory bank:

$$M_b[\tilde{I}] \leftarrow \{\mathbf{z}_i\}_{i=1}^{|\tilde{I}|}. \quad (5)$$

where \tilde{I} is the corresponding index set of this batch of samples in M_b . After *epoch1* training set training, the model can learn a better relation representation.

3.5 Selecting Typical Samples for Memory

In order to make the model not forget the relevant knowledge of the old task when it learns the new task, some samples need to be stored in M_r . Inspired by (Han et al., 2020; Cui et al., 2021), we use k-means to cluster each relation, where the number of clusters is the number of samples that need to be stored for each class. Then, the relation representation closest to the center is selected and stored in memory for each cluster.

3.6 Consistent Representation Learning

After learning a new task, the representation of the old relation in the space may change. In order to make the encoder not change the knowledge of the old task while learning the new task, we propose two replay strategies to learn consistent representation for alleviating this problem: contrastive replay and knowledge distillation. Figure 1 shows the main flow of consistent representation learning.

Contrastive Replay with Memory Bank After the new task learning is over, we use the new task to train the encoder to further train the encoder by replaying the samples stored in memory M_k . After the learning of the current task is over, we use the same method in Section 3.4 to replay the samples stored in memory M_k .

The difference here is that each batch uses all the samples in the entire memory bank for contrastive learning, as follows:

$$\mathcal{L}_{CR} = \sum_{i \in I} \frac{-1}{|P(i)|} \sum_{p \in P(i)} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{j \in \tilde{S}_I} \exp(z_i \cdot z_j / \tau)}, \quad (6)$$

where \tilde{S}_I represents the set of indices of all samples in \tilde{M}_b . \tilde{M}_b is the memory bank, which stores the normalized representation of all samples in M_k .

By replaying the samples in memory, the encoder can alleviate the forgetting of previously learned knowledge, and at the same time, consolidate the knowledge learned in the current task. However, contrastive replay allows the encoder to train on a small number of samples, which risks overfitting. On the other hand, it may change the distribution of relations in the previous task. Therefore, we propose knowledge distillation to make up for this shortcoming.

Knowledge Distillation for Relieve Forgetting

We hope that the model can retain the semantic knowledge between relations in historical tasks.

Therefore, before the encoder is trained on a task, we use the similarity metric between the relations in memory as **Memory Knowledge**. Then use the knowledge distillation to relieve the model from forgetting this knowledge.

Specifically, the samples in the memory are encoded first, and then the prototype of each class is calculated:

$$p_c = \sum_{i=1}^O z_i^c, \quad (7)$$

where O is the number of memory size, z_i^c is the relation representation belonging to class c . Then, the cosine similarity between the classes is calculated to represent the knowledge learned in the memory:

$$a_{ij} = \frac{p_i^T p_j}{\|p_i\| \|p_j\|}, \quad (8)$$

where a_{ij} is the cosine similarity between prototype i and j .

When performing memory replay, we use KL divergence to make the encoder retain the knowledge of the old task.

$$\mathcal{L}_{KL} = \sum_i KL(P_i || Q_i), \quad (9)$$

where $P_i = \{p_{ij}\}_{j=1}^{|\hat{R}_k|}$ is the metric distribution of the prototype before training, and $p_{ij} = \frac{\exp(a_{ij}/\tau)}{\sum_j \exp(a_{ij}/\tau)}$. Similarly, $Q_i = \{q_{ij}\}_{j=1}^{|\hat{R}_k|}$ is the metric distribution of calculate the temporary prototype from the memory bank during training, and $q_{ij} = \frac{\exp(\tilde{a}_{ij}/\tau)}{\sum_j \exp(\tilde{a}_{ij}/\tau)}$. \tilde{a} is the **Embedding Knowledge** of the memory M_k , which is the cosine similarity between temporary prototypes. The temporary prototype is dynamically calculated in each batch based on the memory bank \tilde{M}_b .

3.7 NCM for Prediction

To predict a label for a test sample x , the nearest class mean (NCM) (Mai et al., 2021) compares the embedding of x with all the prototypes of memory and assigns the class label with the most similar prototype:

$$p_c = \frac{1}{n_c} \sum_i \mathbf{E}(\bar{x}_i) \cdot \mathcal{K}\{y_i = c\}, \quad (10)$$

$$y^* = \underset{c=1, \dots, k}{\operatorname{argmin}} \|f(x) - p_c\|,$$

where $\bar{x} \in M_k$ is stored sample, and y^* is a predicted label. Since the NCM classifier compares the embedding of the test sample with prototypes, it

FewRel										
Model	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
EA-EMR	89.0	69.0	59.1	54.2	47.8	46.1	43.1	40.7	38.6	35.2
EMAR	88.5	73.2	66.6	63.8	55.8	54.3	52.9	50.9	48.8	46.3
CML	91.2	74.8	68.2	58.2	53.7	50.4	47.8	44.4	43.1	39.7
EMAR+BERT	98.8	89.1	89.5	85.7	83.6	84.8	79.3	80.0	77.1	73.8
RP-CRE	97.9	92.7	91.6	89.2	88.4	86.8	85.1	84.1	82.2	81.5
RP-CRE [†]	98.4	95.2	93.1	91.4	90.8	88.8	87.6	86.8	85.2	83.9
CRL	98.3	95.4	93.4	92.0	91.0	89.7	88.3	87.0	85.6	84.4
w/o KL	98.3	95.2	93.1	91.5	90.4	89.0	87.7	86.3	84.9	83.4
w/o CR	98.3	94.8	92.2	90.7	89.4	87.6	86.5	85.0	83.7	82.0

TACRED										
Model	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10
EA-EMR	47.5	40.1	38.3	29.9	24	27.3	26.9	25.8	22.9	19.8
EMAR	73.6	57.0	48.3	42.3	37.7	34.0	32.6	30.0	27.6	25.1
CML	57.2	51.4	41.3	39.3	35.9	28.9	27.3	26.9	24.8	23.4
EMAR+BERT	96.6	85.7	81	78.6	73.9	72.3	71.7	72.2	72.6	71.0
RP-CRE	97.6	90.6	86.1	82.4	79.8	77.2	75.1	73.7	72.4	72.4
RP-CRE [†]	97.8	92.3	91.0	87.3	84.2	82.7	79.8	78.8	78.6	77.3
CRL	98.1	94.7	91.6	87.0	86.3	84.5	82.9	81.8	81.8	80.7
w/o KL	98.1	94.2	91.7	87.1	86.6	84.4	82.2	81.5	81.0	80.1
w/o CR	98.1	93.2	90.1	85.8	83.2	81.2	79.4	77.4	76.8	75.9

Table 1: Accuracy (%) on all observed relations (which will continue to accumulate over time) at the stage of learning current task. The method marked by [†] represents the results generated from open source code¹ and the other baseline results copied from the original paper (Cui et al., 2021)

does not require an additional FC layer. Therefore, new classes can be added without any architecture modification.

4 Experiments

4.1 Datasets

Our experiments are conducted on two benchmark datasets: in the experiment, the training-test-validation that the split ratio is 3:1:1.

FewRel (Han et al., 2018) It is a RE dataset that contains 80 relations, each with 700 instances. Following the experimental settings by Wang et al. (2019), the original train and valid set of FewRel are used for experimental, which contains 80 classes.

TACRED (Zhang et al., 2017) It is a large-scale RE dataset containing 42 relations (including no relations) and 106,264 samples, built on news networks and online documents. Compared with FewRel, the samples in TACRED are imbalanced. Following Cui et al. (2021), the number of training samples for each relation is limited to 320 and the number of test samples of relation to 40.

¹<https://github.com/fd2014cl/RP-CRE>

4.2 Evaluation Metrics

Average accuracy is a better measure of the effect of catastrophic forgetting because it emphasizes the model’s performance on earlier tasks (Han et al., 2020; Cui et al., 2021). This paper evaluates the model by using the average accuracy of K tasks at each step.

4.3 Baselines

We evaluate CRL and several baselines on benchmarks for comparison:

(1) EA-EMR (Wang et al., 2019) introduced a memory replay and embedding alignment mechanism to maintain memory and alleviate embedding distortion during training for new tasks.

(2) EMAR (Han et al., 2020) constructs a memory activation and reconsolidation mechanism to alleviate the catastrophic forgetting problem in CRE.

(3) CML (Wu et al., 2021) proposed a curriculum-meta learning method to alleviate the order sensitivity and catastrophic forgetting in CRE.

(4) RP-CRE (Cui et al., 2021) achieves enhanced performance by utilizing relation prototypes to refine sample embeddings, thereby effectively avoiding catastrophic forgetting.




Figure 2: Comparison of model’s dependence on memory size, it shows that our model has a light dependence on memory size. The X-axis is the serial ID of the current task, Y-axis is the accuracy of the standard model on the test set from all observed relations at current stage.

4.4 Training Details and Parameters Setting

A completely random sampling strategy at the relation level is adopted. It simulates ten tasks by randomly dividing all relations of the dataset into 10 sets to simulate 10 tasks, as suggested in (Cui et al., 2021). For a fair comparison, we set the random seed of the experiment to be the same as the seed in (Cui et al., 2021), so that the task sequence is exactly the same. Note that our reproduced model RP-CRE † and CRL use strictly the same experimental environment. In order to facilitate the reproduction of our experimental results, the proposed method source code and detailed hyperparameters are provided on Github².

4.5 Results and Discussion

Table 1 shows the results of the proposed methods and baselines ones compared on two datasets, where RP-CRE † is reproduced under the same conditions based on open source code. We also ablated knowledge distillation and contrastive replay for consistent representation learning. CRL (w/o KL) and CRL (w/o CR) respectively refer to removing knowledge distillation loss \mathcal{L}_{KL} and contrastive replay loss \mathcal{L}_{CR} when replaying memory. From the table, some conclusions can be drawn:

(1) Our proposed CRL is significantly better than other baselines and achieves state-of-the-art performance in the vast majority of settings. Compared with RP-CRE, our model also produces apparent advantages. It proves that CRL can learn better consistent relation representations and is more stable in the process of continual learning.

(2) It is observed that all baselines perform worse on the TACRED dataset. The primary reason for this result is that TACRED is an imbalanced dataset. However, our model performs better than

RP-CRE’s last task on TACRED (3.4% higher than RP-CRE), which is more significant than the improvement (0.5%) on the class-balanced dataset FewRel. It shows that our model is more robust to scenarios with class-imbalanced.

(3) Comparing CRL and CRL (w/o KL), not adopting knowledge distillation during training can cause the model to drop 1% and 0.6% on FewRel and TACRED, respectively. The experimental results show that knowledge distillation can uniformly alleviate the model’s forgetting of previous knowledge to learn a better consistent representation.

(4) Comparing CRL and CRL (w/o CR), removing L during memory replay caused the model to drop 2.4% and 4.8% on FewRel and TACRED, respectively. The reason for the significant drop is that only adopting \mathcal{L}_{KL} cannot make the model review the samples of the current task, which leads to overfitting in the historical relations during replay.

4.6 Effect of Memory Size

The memory size is the number of memory samples needed for each relation. In this section, we will study the impact of memory size on the performance of our model and RP-CRE. We compare three memory sizes: 5, 10, and 20. The experimental results are shown in Figure 2.

We choose RP-CRE as the main competitor, where all configurations and task sequence remain unchanged. (1) As the size of the memory decreases, the performance of the model tends to decline, which shows that the size of the memory is a key factor that affects continuous learning and learning. But our model is more stable than RP-CRE (the performance gap in the final task), especially on the TACRED dataset. (2) On both FewRel and TACRED, CRL keeps the best perfor-

²<https://github.com/thuiar/CRL>




Figure 3: A visualization of relation representation learned from task 1 test set by RP-CRE and CRL at different task.

mance under different memory sizes and produces obvious advantages in small memory. It indicates that utilizing consistent representation learning is a more effective way to utilize memory than the existing memory-based CRE method.

4.7 Effect of Consistent Representation Learning

In order to explore the long-term effects of consistency representation learning in continual relation extraction, we tested our model and RP-CRE on TACRED to observe the changes in the embedding space of old tasks as new tasks continue to increase. The model performs feature extraction on all samples in the test set in task 1 at the end of tasks 1, 4, 7, and 10. Then t-SNE is used to represent the dimensionality reduction relation representation. All samples on the test set of task 1 are drawn, where different color points represent different ground-truth labels. The visualization results are shown in Figure 3.

From Figure 3, we can see that although the relation embeddings of RP-CRE are clustered and separated in each class after prototype refinement, as new tasks are continuously learned, the data embedding of task 1 is obviously scattered. In contrast, our model retains a good separation between classes, while the data embedding within classes is compact and has a certain diversity. In addition, we can see that our model has relatively stable changes in the distribution of different classes

in task 1, and retains the knowledge of historical tasks with training. This is mainly because our model learns through supervised comparison, and explicitly emphasizes that the samples in historical memory are compact within the class and far away from each other. And the knowledge of historical memory is preserved through the distillation of memory knowledge. Because knowledge distillation preserves the distance distribution between classes, it can make up for the contrastive learning to over-optimize the distance between classes to prevent overfitting.

5 Conclusions and Future Work

This paper proposes a novel consistent representation learning method for the CRE task, mainly through contrastive learning and knowledge distillation when replaying memory. Specifically, we use supervised contrastive learning based on a memory bank to train each new task so that the model can effectively learn the feature representation. In addition, in order to prevent the catastrophic forgetting of the old task, we contrast and replay the memory samples, and at the same time, make the model retain the knowledge of the relation between the historical tasks through the knowledge distillation. Our method can better learn consistent representations to alleviate catastrophic forgetting effectively. Extensive experiments on two benchmark data sets show that our method significantly improves the performance of the most advanced technology and

demonstrates powerful representation learning capabilities. In the future, we will continue to study cross-domain continual relation extraction to acquire ever-increasing knowledge.

Acknowledgments

This paper is funded by National Natural Science Foundation of China (Grant No. 62173195) and Natural Science Foundation of Hebei Province, China (pre-research No. F2022208006). We would like to thank the anonymous reviewers for their valuable feedback.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2018. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Tianqi Chen, Ian Goodfellow, and Jonathon Shlens. 2015. Net2net: Accelerating learning via knowledge transfer. *arXiv preprint arXiv:1511.05641*.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Xinlei Chen and Kaiming He. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758.
- Li Cui, Deqing Yang, Jiaxin Yu, Chengwei Hu, Jiayang Cheng, Jingjie Yi, and Yanghua Xiao. 2021. Refining sample embeddings with relation prototypes to enhance continual relation extraction. In *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)*, pages 232–243.
- Cyprien de Masson D’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. 2019. Episodic memory in lifelong language learning. *Advances in Neural Information Processing Systems*, 32.
- 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-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Songlin Dong, Xiaopeng Hong, Xiaoyu Tao, Xinyuan Chang, Xing Wei, and Yihong Gong. 2021. Few-shot class-incremental learning via relation knowledge distillation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 1255–1263.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. 2017. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*.
- Xu Han, Yi Dai, Tianyu Gao, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2020. Continual relation learning via episodic memory activation and reconsolidation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6429–6440.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. **FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Dan Hendrycks and Thomas Dietterich. 2019. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*.
- Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. 2021. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526.

- Junnan Li, Pan Zhou, Caiming Xiong, and Steven CH Hoi. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Hao Liu and Pieter Abbeel. 2020. Hybrid discriminative-generative training via contrastive learning. *arXiv preprint arXiv:2007.09070*.
- Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. 2018. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 2262–2268. IEEE.
- David Lopez-Paz and Marc Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476.
- Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner. 2021. Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3589–3599.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019a. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of ACL*.
- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019b. [Matching the blanks: Distributional similarity for relation learning](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2895–2905. Association for Computational Linguistics.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. Lamol: Language modeling for lifelong language learning. *arXiv preprint arXiv:1909.03329*.
- Chongyang Tao, Shen Gao, Mingyue Shang, Wei Wu, Dongyan Zhao, and Rui Yan. 2018. [Get the point of my utterance! learning towards effective responses with multi-head attention mechanism](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4418–4424. ijcai.org.
- Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019. Sentence embedding alignment for lifelong relation extraction. In *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)*, pages 796–806.
- Tongtong Wu, Xuekai Li, Yuan-Fang Li, Reza Hafari, Guilin Qi, Yujin Zhu, and Guoqiang Xu. 2021. Curriculum-meta learning for order-robust continual relation extraction. *CoRR, abs/2101.01926*.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.
- Chenyan Xiong, Russell Power, and Jamie Callan. 2017. [Explicit semantic ranking for academic search via knowledge graph embedding](#). In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, pages 1271–1279. ACM.
- Shipeng Yan, Jiangwei Xie, and Xuming He. 2021. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-based bidirectional long short-term memory networks for relation classification. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 2: Short papers)*, pages 207–212.