

Curriculum-Meta Learning for Order-Robust Continual Relation Extraction

Tongtong Wu^{1*}, Xuekai Li¹, Yuan-Fang Li², Gholamreza Haffari², Guilin Qi^{1†},
Yujin Zhu³, Guoqiang Xu³

¹School of Computer Science and Engineering, Southeast University, Nanjing, China

²Faculty of Information Technology, Monash University, Melbourne, Australia

³Gamma Lab, Ping An OneConnect, Shanghai, China

{wutong8023,lxk,gqi}@seu.edu.cn, {yuanfang.li,Gholamreza.Haffari}@monash.edu,
{zhuyujin204, xuguoqiang371}@ocft.com

Abstract

Continual relation extraction is an important task that focuses on extracting new facts incrementally from unstructured text. Given the sequential arrival order of the relations, this task is prone to two serious challenges, namely catastrophic forgetting and order-sensitivity. We propose a novel curriculum-meta learning method to tackle the above two challenges in continual relation extraction. We combine meta learning and curriculum learning to quickly adapt model parameters to a new task and to reduce interference of previously seen tasks on the current task. We design a novel relation representation learning method through the distribution of domain and range types of relations. Such representations are utilized to quantify the difficulty of tasks for the construction of curricula. Moreover, we also present novel difficulty-based metrics to quantitatively measure the extent of order-sensitivity of a given model, suggesting new ways to evaluate model robustness. Our comprehensive experiments on three benchmark datasets show that our proposed method outperforms the state-of-the-art techniques. The code is available at <https://github.com/wutong8023/AAAI-CML>.

Introduction

Relation extraction (?) aims at extracting structured facts as triples from unstructured text. As an essential component of information extraction, relation extraction has been widely utilized in downstream applications such as knowledge base construction (?) and population (?). However, given the continuous and iterative nature of the update process, continual relation extraction (??) is a more realistic and useful setting. Yet due to the limitations of storage and computational resources, it is impractical to grant the relation extractor access to all the training instances in previously seen tasks. Thus, this *continual learning* formulation is in contrast to the conventional relation extraction setting where the extractor is generally trained from scratch with the full access to the training corpus.

*This work has been done by Tongtong Wu during the visiting period at Monash University, and the original idea was generated in the internship at Gamma Lab, Ping An OneConnect.

†Contact author.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Catastrophic Forgetting (CF) is a well-known problem in continual learning (?). The problem is that when a neural network is utilized to learn a sequence of tasks, the learning of the later tasks may degrade the performance of the learned model for the previous tasks. Various recent works tackle the CF problem, including consolidation-based methods (??), dynamic architecture methods (??), or memory-based methods (??). These methods have been demonstrated in simple image classification tasks. Yet the memory-based methods have been proven to be the most promising for NLP applications. EA-EMR (?) proposes a sentence embedding alignment mechanism in memory maintenance and adopt it to continual relation extraction learning. Based on EA-EMR, MLLRE (?) introduces a meta-learning framework for fast adaptation and EMAR (?) introduces a multi-turn co-training procedure for memory-consolidation. Most of these methods explore CF problem in the overall performance of task sequences, but they lack the insight analysis of the characteristics of each subtask and the corresponding model performance.

Order-sensitivity (OS) is another major problem in continual learning, which is relatively under-explored (?). It refers to the phenomenon that the performance of tasks varies based on the order of the task arrival sequence. This is due to not only the CF incurred by the different sequences of previous tasks but also the unidirectional knowledge transfer from the previous tasks. Order-sensitivity can be problematic in various aspects: (i) ethical AI considerations in continual learning, e.g. fairness in the medical domain (?); (ii) bench-marking of continual learning algorithms as most of the existing works pick an arbitrary and random sequence of the given tasks for evaluation (?); (iii) uncertainty to the quality of extracted knowledge in the realistic scenario for knowledge base population, where the model is faced with only one sequence.

In this paper, we introduce the curriculum-meta learning (CML) method to tackle both the catastrophic forgetting and order-sensitivity problems. Taking a memory-based approach, CML is based on the following observations about the catastrophic forgetting and order sensitivity issues of the previous works: (i) over-fitting to the experience memory, indicating that the performance on any task will decrease as

training progresses, and (ii) the interference between similar tasks, indicating that the model performs better on less intrusive tasks. We therefore design a mechanism which selectively reduces the replay frequency of memory to avoid over-fitting, and steer the model to learn the bias between the current task and previous most similar tasks to reduce the order-sensitivity.

Our CML method contains two steps. In the first step, it samples instances from the memory based on the difficulty of the previous tasks for the current task, resulting in a curriculum for continual learning. Then, it trains the model on both the curriculum and training instances of the current task. We further introduce a knowledge-based method to quantify task difficulty according to the similarity of pairs of relations. Taking a relation as a function mapping to named entities in its domain and range, we define a similarity measure between two relations based on the conceptual distribution of their head and tail entities. Our contributions are summarized as follows:

- We propose a novel curriculum-meta learning method to tackle the order-sensitivity and catastrophic forgetting problems in continual relation extraction.
- We introduce a new relation representation learning method via the conceptual distribution of head and tail entities of relations, which is utilized to quantify the difficulty of each relation extraction task for constructing the curriculum.
- We conduct comprehensive experiments to analyze the order-sensitivity and catastrophic forgetting problems in state-of-the-art methods, and empirically demonstrate that our proposed method outperforms the state-of-the-art methods on three benchmark datasets.

Related Work

The conventional relation extraction methods could be categorized into three domains by the way data is used: supervised methods (????), semi-supervised methods (??), and distantly supervised methods (??). Most of these methods assume a predefined relation schema and thus cannot be easily generalized to new relations.

In this paper, we address the continual relation learning problem (?), a relatively new and less investigated task. Continual learning in general faces two major challenges: catastrophic forgetting and order-sensitivity.

Catastrophic forgetting (CF) is a prominent line of research in continual learning (??). Methods addressing CF can be broadly divided into three categories. (i) Consolidation-based methods (????) consolidate model parameters important to previous tasks and reduce their learning weights. These methods employ sophisticated mechanisms to evaluate parameter importance for tasks. (ii) Dynamic architecture methods (??) dynamically expand model architectures to learn new tasks and effectively prevent forgetting of old tasks. Sizes of these methods grow dramatically with new tasks, making them unsuitable for NLP applications. (iii) Memory-based methods (????) remember a few examples in old tasks and continually learn them

with emerging new tasks to alleviate catastrophic forgetting. Among these methods, the memory-based methods have been proven to be the most promising for NLP applications (??), including both relation learning (??).

Order-sensitivity (OS) (??) is another major problem in continual learning that is relatively under-explored. It is the phenomenon that a model’s performance is sensitive to the order in which tasks arrive. In this paper, we tackle this problem by leveraging a curriculum learning method (?). Briefly, we construct our curriculum by the similarity of tasks, thus minimizing the impact and interference of previous tasks.

Curriculum-Meta Learning

Problem Formulation. In continual relation extraction, given a sequence of K tasks $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_K\}$, each task \mathcal{T}_k is a conventional supervised classification task, containing a series of examples and their corresponding labels $\{(x^{(i)}, y^{(i)})\}$, where $x^{(i)}$ is the input data, containing the natural-language context and the candidate relations, and $y^{(i)}$ is the ground-truth relation label of the context. The model $f_\theta(\cdot)$ can access the training data of the current task \mathcal{D}_k^{train} and is trained by optimising a loss function $l(f_\theta(x), y)$. The goal of continual learning is to train the model $f_\theta(\cdot)$ such that it continually learns new tasks while avoiding catastrophically forgetting the previously learned tasks. Due to various constraints, the learner is typically allowed to maintain and observe only a subset of the training data of the previous tasks, contained in a memory set \mathcal{M} .

The performance of the model $f_\theta(\cdot)$ is measured in the conventional way, by *whole accuracy* $Acc_w = acc_{f, \mathcal{D}^{test}}$, on the entire test set, where $\mathcal{D}^{test} = \bigcup_{i=1}^K \mathcal{D}_i^{test}$. Moreover, model performance at task k is evaluated with *average accuracy* on the test sets of all the tasks up to this task in the sequence $Acc_a = \frac{1}{k} \sum_{i=1}^k acc_{f,i}$. Average accuracy is a better measure of the effect of catastrophic forgetting as it emphasizes on a model’s performance of earlier tasks.

Framework. Our curriculum-meta learning (CML) framework is described in Algorithm ?. CML maintains initialization parameters θ_t and a memory set \mathcal{M} that stores the prototype instances of previous tasks. It performs the following operations at each time step t during the learning phase. (1) The meta-learner L fetches the initialization parameters θ_t from the memory to initialize the model $f_{\theta_t}(\cdot)$. (2) L replays on the curriculum set $D_{t,i}^{curri}$ which is sampled and sorted by the knowledge-based curriculum module. (3) L trains on the support set $D_{t,i}^{train}$ of the current task \mathcal{T}_t . (4) Finally, L updates the learned parameters θ_{t+1} and stores a small number of prototype instances of the current task into the memory. During the evaluation phase, the trained model is given a target set with labeled unseen instances from all observed tasks (See Appendix A for the workflow of CML.)

We will introduce the framework in terms of (1) the utilization of the initialization parameters (i.e. meta training) and (2) the utilization of the memory set (i.e. the curriculum-based memory replay).

Meta Training. Meta learning, or learning to learn, aims

Algorithm 1: Curriculum-Meta Learning

Input: Stream of incoming tasks $\mathcal{T}_1, \mathcal{T}_2, \dots$; Classification model f_θ , Step size ϵ , Learning rate α , Relations embedding r , Memory buffer \mathcal{M} , Curriculum size k , Curriculum instance size n , Knowledge-based curriculum module g_ϕ

```
1 Initialize  $\theta$ 
2 while there are still tasking do
3   Retrieve current task  $\mathcal{T}_t$ 
4   Initialize  $\theta_t \leftarrow \theta$ 
5   while not convergence do
6     if  $\mathcal{M}$  is not empty then
7       Set  $\mathcal{D}_t^{train} \leftarrow \mathcal{D}_t^{train} \cup \mathcal{M}$ 
8     for each relation  $R_i$  in  $\mathcal{T}_t$  do
9       Initialize  $\theta_{t,i} \leftarrow \theta_t$ 
10      Sample  $D_{t,i}^{train}$  from  $D_t^{train}$  for  $R_i$ 
11      if  $\mathcal{T}_t$  is not the first task then
12        Sample  $k$  curriculum relations from  $\mathcal{M}$ 
13        with  $g_\phi$ .
14        Sample  $n$  instances for each curriculum
15        relation
16        Construct sorted mini-batches, each
17        contains  $D_{t,i}^{curri}$  of  $k \times n$  instances.
18        Evaluate  $\nabla_{\theta_{t,i}} L_{R_i}(f_{\theta_{t,i}})$  using  $D_{t,i}^{curri}$ 
19        Update  $\theta'_{t,i} = Adam(\theta_{t,i}, L_{R_i}(f_{\theta_{t,i}}), \alpha)$ 
20      Evaluate  $\nabla_{\theta'_{t,i}} L_{R_i}(f_{\theta'_{t,i}})$  using  $D_{t,i}^{train}$ 
21      Update  $\theta_{t,i}^* = Adam(\theta'_{t,i}, L_{R_i}(f_{\theta'_{t,i}}), \alpha)$ 
22    Update  $\theta_{t+1} = \theta_t + \frac{\epsilon}{N} \sum_{i=1}^N (\theta_{t,i}^* - \theta_t)$ 
23    Sample  $sub\mathcal{T}_t$  from  $\mathcal{T}_t$ 
24    Update  $\mathcal{M} \leftarrow \mathcal{M} \cup sub\mathcal{T}_t$ 
25    Update  $\theta \leftarrow \theta_t$ 
26    Fine-tune  $\theta$  on  $\mathcal{M}$ 
```

at developing algorithms to learn the generic knowledge of how to solve tasks from a given distribution of tasks. With a given basic relation extraction model (?) $f_\theta(\cdot)$ parameterized by θ , we employ the gradient-based meta-learning method (?) to learn a prior initialization θ_t at each time step t . During adaption to a new task, the model parameters θ are quickly updated from θ_t to the task-specific θ_t^* with a few steps of gradient descent. Formally, the meta learner L updates $f_\theta(\cdot)$ that is optimized for the following objectives:

$$\min_{\theta} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} [\mathcal{L}(\theta^*)] = \min_{\theta} \mathbb{E}_{\mathcal{T} \sim p(\mathcal{T})} [\mathcal{L}(\mathcal{U}(\mathcal{D}_{\mathcal{T}}\theta))] \quad (1)$$

where $\mathcal{D}_{\mathcal{T}}$ is the training data, \mathcal{L} is the loss function for task \mathcal{T} , and \mathcal{U} is the optimizer of $f_\theta(\cdot)$. Then, when it converges on the current task, the model will generate the initialization parameter θ_{t+1} for the next time step $t + 1$:

$$\theta_{t+1} = \theta + \frac{\epsilon}{n} \sum_{t=1}^n (\theta_t^* - \theta_t), \quad (2)$$

where the θ_t^* is the updated parameter for the current task \mathcal{T}_k at time step t , and n is the number of instances which may be processed in parallel at a time step.

Curriculum-based Memory Replay. Meta learner L reviews the previous tasks in an orderly way before learning

the new task. Here, we denote by $g_\phi(\cdot)$ a function to represent the teacher which prepares the curriculum for the student network (i.e. the relation extractor $f_\theta(\cdot)$) for replay. Different from conventional experience-replay based models, the teacher function needs to master three skills:

1. Assessing the difficulty of tasks. When a new task arrives, this function calculates which of all observed previous tasks interferes with the current task.
2. Sampling instances from the memory. By sampling, we can reduce the time consumption in the replay stage and alleviate the over-fitting problems caused by the high frequency of updates on the memory.
3. Ranking the sampled instances by a certain strategy. The teacher instructs the student model to learn the bias between the current task and observed similar tasks in the most efficient way.

We sample the memory randomly and sort the sampled instances according to the difficulty of each previous task with respect to the current task. Based on the above requirements, we implement a knowledge-based curriculum module, which is introduced in the next section.

Knowledge-based Curriculum

Intrinsically, order-sensitivity is caused by a model’s inability to guarantee optimal performance for all previous tasks. However, from the perspective of experiments, order-sensitivity is closely related to the unbalanced forgetting rate (or the unbalanced difficulty) of different tasks, where we assume that task difficulty is due to the interactions between semantically similar relations in an observed task sequence. Intuitively, if the conceptual distribution of two relations is similar, these two relations tend to be expressed in similar natural language contexts, such as the relations “father” and “mother”.

Semantic Embedding-based Difficulty Function. To formalize this intuition, we define a difficulty estimation function based on the semantic embeddings of relations in each task. Given a set of K tasks $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_K\}$, the difficulty of task \mathcal{T}_i is defined as:

$$Dl_i := \frac{1}{K-1} \sum_{\substack{j=1 \\ j \neq i}}^K S_i^j \quad (3)$$

S_i^j is the similarity score between tasks \mathcal{T}_i and \mathcal{T}_j , which is defined as the average similarity among relation pairs from the two tasks:

$$S_i^j := \frac{1}{M \times N} \sum_{m=1}^M \sum_{n=1}^N s_m^n, \quad (4)$$

where M and N are the numbers of relations in each task respectively, and s_m^n calculates the Cosine similarity between the embeddings of the two relations: $s_m^n := \cos(emd_m, emd_n)$. Using s_m^n , we calculate the difficulty of each relation in the memory with respect to the relations in the current task, in order to sort and sample the relations stored in memory into the final curriculum.

Relation Representation Learning. In order to calculate the semantic embedding of each relation, inspired by (?), we introduce a knowledge- and distribution-based representation learning method. Intuitively, the representation of a relation is learned from the types of its head and tail entities. Consider a knowledge graph $\mathcal{G} = \{(h, r, t) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}\}$, where h and t are the head and tail entities, r is the relation between them, and \mathcal{E} and \mathcal{R} represent the sets of entities and relations respectively. We reduce the relation representation learning task into the problem of learning the conceptual distribution of each relation which is optimized based on the following objective:

$$\begin{aligned} \min_{\phi} \mathcal{L}(\phi; \mathcal{G}) = \\ \min_{\phi} \sum_{(h', t'; r) \in \mathcal{G}} [-\log P_{\phi}(h'|r) - \log P_{\phi}(t'|r)] \end{aligned} \quad (5)$$

where h' and t' are the concepts (i.e. the hypernyms obtained from the knowledge graph) of the head and tail entities respectively. $P_{\phi}(h'|r) = \exp(NN_{\phi_1}(h', r))$, and $P_{\phi}(t'|r) = \exp(NN_{\phi_2}(t', r))$, where $NN_{\phi}(a, b) = MLP_{\phi}(\mathbf{a})^{\top} \mathbf{b}$ is two-layer neural network parameterized with ϕ . Finally, we obtain two representations emd_r^h and emd_r^t for each relation, which indicate the conceptual-distribution of head entities and tail entities respectively. We concatenate these two embeddings to generate the final representation of the relation $emd_r := [emd_r^h; emd_r^t]$.

Experiments

In this section, we aim to empirically address the following research questions related to our contributions:

RQ1: Why and to what extent do current memory replay-based approaches suffer from catastrophic forgetting and order-sensitivity?

RQ2: How to qualitatively and quantitatively understand task difficulty?

RQ3: Compared with the state-of-the-art methods, can our method (curriculum-meta learning with the knowledge-based curriculum) effectively alleviate catastrophic forgetting and order-sensitivity?

Datasets. We conduct our experiments on three datasets, including Continual-FewRel, Continual-SimpleQuestions, and Continual-TACRED, which were introduced in (?). FewRel (?) is a labelled dataset which contains 80 relations and 700 instances per relation. SimpleQuestions is a knowledge-based question answering dataset containing single-relation questions (?), from which a relation extraction dataset was extracted (?). The relation extraction dataset contains 1,785 relations and 72,238 training instances. TACRED (?) is a well-constructed RE dataset that contains 42 relations and 21,784 examples. Considering the special relation “n/a” (i.e. not available) in TACRED, we follow (?) and filter out these examples with the relation “n/a” and use the remaining 13,012 examples for Continual-TACRED.

Following ??, we partition the relations of each dataset into some groups and then consider each group of relations as a distinct task T_k . We form training and testing set for

each task, based on the instances in the original dataset labeled by the relations in the task. Following the previous work, we employ two relation partitioning methods. Firstly, the unbalanced division is based on *clustering*, using the averaged word embeddings (?) of relation names with the K-means clustering algorithm (?). Secondly, the random partitioning into groups with a similar number of relations (?). For Continual-FewRel, we partition its 80 relations into 10 distinct tasks. Similarly, we partition the 1,785 relations in Continual-SimpleQuestions into 20 disjoint tasks, so as well as partition the 41 relations in Continual-TACRED into 10 tasks.

Evaluation Metrics. We employ the following four metrics to measure model performance. Note that the last two metrics, the average forgetting rate and the error bound, are new metrics we propose in this paper.

Whole Accuracy of the resulting model at the end of the continual learning process on the full test sets of all tasks, $Acc_w := acc_{f, \mathcal{D}_{test}}$.

Average Accuracy of the resulting model trained on task T_k on all the test sets of all tasks seen up to stage k of the continual learning process, $Acc_a := \frac{1}{k} \sum_{i=1}^k acc_{f,i}$. Compared to Acc_w , Acc_a highlights the catastrophic forgetting problem. However, as we will empirically show, Acc_a is subject to *order-sensitivity* of the tasks sequence, and thus does not accurately measure the level of forgetting on a specific task.

Average Forgetting Rate for task j after k time steps, Fr_{avg}^j is a new metric to evaluate task-specific model performance on order-sensitivity.

$$Fr_{avg}^j := \frac{1}{k-1} \sum_{i=1}^{k-1} \frac{\overline{acc}_{i+1}^j - \overline{acc}_i^j}{\overline{acc}_i^j}, \quad (6)$$

where \overline{acc}_i^j is the model’s average performance on a specific task T_j when it appears in the i th position of distinct task permutations:

$$\overline{acc}_i^j := \frac{1}{(J-1)!} \sum_{\pi \in \Pi_{[1, \dots, J]} \text{ st } \pi_i = j} acc_i(\pi) \quad (7)$$

where $acc_i(\pi)$ is the final accuracy on task T_i of the model trained on the permutation π , $\Pi_{[1, \dots, J]}$ is the set of all permutations of the tasks $\{T_1, \dots, T_J\}$, and π_i is the index of Task T_i of the sequence. We note that the number of all permutations in which task T_i is fixed at position j is $(J-1)!$. Of course we may not be able to exactly compute $\overline{acc}_{f,i}^j$ as the size of possible tasks permutations grows exponentially. Therefore, we estimate this quantity by Monte Carlo sampling of some permutations.

Error Bound is a new metric to evaluate the overall model performance regarding order-sensitivity.

$$EB := Z_{\frac{\delta}{2}} \times \frac{\delta}{\sqrt{n}}, \quad (8)$$

where $Z_{\frac{\delta}{2}}$ is the confidence coefficient of confidence level α , and δ is the standard deviation of accuracy obtained

from n distinct task permutations. Note that a model with a lower error bound shows better robustness and less order-sensitivity for the input sequences.

Baseline Models.

We compare our proposed CML with knowledge-based curriculum-meta learning with the following baseline models, among which Vanilla is employed as the base learner for both CML and the other models (see Appendix B for hyper-parameters.):

1. Vanilla (?), which is the basic model for conventional supervised relation extraction not specifically designed for the continual learning setup.
2. EWC (?), which adopts elastic weight consolidation to add special L_2 regularization on parameter changes. Then, EWC uses Fisher information to measure the parameter importance to old tasks, and slow down the update of those parameters important to old tasks.
3. AGEM (?), which takes the gradient on sampled memorized examples from memory as the only constraint on the optimization directions of the current task.
4. EA-EMR (?), which maintains a memory of previous tasks to alleviate the catastrophic forgetting problem.
5. MLLRE (?), which leverages meta-learning to improve the usage efficiency of training instances.
6. EMAR (?), which introduces episodic memory activation and reconsolidation to continual relation learning.

Main Results

To evaluate the overall performance of our model CML (**RQ3**), we conduct experiments on the three datasets under both task division methods: unbalanced cluster-based task division and the uniform random task division.

The following observations can be made from Table ??.

- (i) Our model CML achieves the best Acc_w and Acc_a in both settings and on the three datasets in the majority of cases.
- (ii) Specifically, CML achieves the best Acc_w and Acc_a in the two larger datasets Continual-FewRel and Continual-SimpQ.
- (iii) CML obtains the lowest error bounds EB in the majority of cases, demonstrating better stability and lower order-sensitivity.
- (iv) The two task division methods produce the most prominent Acc_a differences on Continual-FewRel for CML. when the data is evenly distributed (i.e. Random), CML’s Acc_a is significantly reduced to be almost equal to Acc_w (from 76.0 to 63.0). On the other two datasets, the performance difference is much less noticeable.

Although the three metrics Acc_a , Acc_w , and EB are good measures of the overall model performance, they do not provide task-specific insights, which we will further discuss in the following subsection.

Analysis of Unbalanced Forgetting

We designed another experiment to better understand the reason for catastrophic forgetting and order-sensitivity

¹<https://github.com/thunlp/ContinualRE>

²<https://github.com/hongwang600/>

(**RQ1**). In this experiment, each task is assigned a fixed ID. Starting with an initial “run” of tasks $0, 1, \dots, 9$, we test model performance on ten different runs generated by the cyclic shift of the initial run. The results of EA-EMR on Continual-Fewrel are summarised in Table ?? (See Appendix C for the result for EMAR, MLLRE, and CML.)

As shown in Table ??, comparing the results in the rows of P_0 and P_9 of each task, we can observe that most tasks see a significant drop in accuracy from P_9 (when the task is the last one seen by the model) to P_0 (when the task is the first one seen by the model), indicating that they suffer from catastrophic forgetting. Intuitively, forgetting on a task reflects an increase in empirical error. We hypothesized that this is most likely due to frequent replay of limited task-related memory, in other words, over-fitting on memory. We tested this hypothesis by adjusting the ratio of training data to memory. Table ?? shows that when memory size is fixed (e.g. 50), more training data (i.e. all vs 200 vs 100) results in poorer performance for all models except EMAR, indicating the model is over-fitting on memory. This is due to a higher replay frequency of the memory. When we fix the ratio of the training data to memory (e.g. 100:25 and 200:50), model performance conforms to the general rule of better performance with more data. Thus, the results shown in Table ?? support our hypothesis.

When reading each column in Table ?? separately, we find that the forgetting rate of each task is different, which cannot be solely explained by the issue of over-fitting on memory. For example, the model performance on T_4 and T_0 is exactly the same (100%) when they appear at the last position P_9 . However, performance on T_4 does not degrade with the advance of its position. On the other hand, we can observe a decreasing trend of performance on T_0 as its position moves back. Moreover, we observe that order-sensitivity may be related to the difficulty of the earlier task in a run, where difficulty refers to a task’s tendency of being more easily forgotten by the model. For instance, we can observe that the final Acc_a of the run ID = 4 experiment (in dark-orange) is the highest, whereas the final Acc_a of the run ID = 6 experiment (in light orange) is the lowest. Also, among all tasks at position P_0 , task T_4 has the highest accuracy while task T_6 has the lowest. Both of the above observations may be explained by task difficulty, which we further study in the next subsection.

Analysis of Task Difficulty

In this section, we present the qualitative and quantitative analyses in order to better understand the difficulty of tasks (**RQ2**). We choose EA-EMR as the case study on the Continual-FewRel dataset with the tasks constructed through clustering.

Qualitative Analysis. Figure ?? shows the t-SNE visualization of the relations, where nodes represent relations, colors represent the tasks, and the distance is calculated from the hidden layer in EA-EMR.

As can be seen from the figure, tasks T_4 and T_8 have only one relation each, and that the relation in T_4 is far away from the others but the relation T_8 is much closer. The differ-

		Continual-FewRel				Continual-SimpQ				Continual-TACRED			
		Acc _w		Acc _a		Acc _w		Acc _a		Acc _w		Acc _a	
Setting	Model	Acc	EB										
Cluster	Vanilla [‡]	16.3	±4.10	19.7	±3.90	60.3	±2.52	58.3	±2.30	12.0	±3.21	8.7	±2.35
	EWC [†]	27.1	±2.32	30.2	±2.10	67.2	±3.16	59.0	±2.20	14.5	±2.51	14.5	±2.90
	AGEM [†]	36.1	±2.51	42.5	±2.63	77.6	±2.11	72.2	±2.72	12.5	±2.24	16.5	±2.20
	EA-EMR [‡]	59.8	±1.50	74.8	±1.30	82.7	±0.48	86.2	±0.33	17.8	±1.01	25.4	±1.17
	EMAR [†]	53.8	±1.30	68.6	±0.71	80.0	±0.83	76.9	±1.39	42.7	±2.92	52.5	±1.74
	MLLRE	56.8	±1.30	70.2	±0.93	84.5	±0.35	86.7	±0.46	34.4	± 0.49	41.2	±1.37
	CML (ours)	60.2	± 0.71	76.0	± 0.24	85.6	± 0.34	87.5	± 0.32	44.4	±1.16	49.3	± 1.01
Random	Vanilla [‡]	19.1	±1.20	19.3	±1.30	55.0	±1.30	55.2	±1.30	10.2	±2.02	10.4	±2.31
	EWC [†]	30.1	±1.07	30.2	±1.05	66.4	±0.81	66.7	±0.83	15.3	±1.70	15.4	±1.79
	AGEM [†]	36.9	±0.80	37.0	±0.83	76.4	±1.02	76.7	±1.01	13.4	±1.47	14.3	±1.62
	EA-EMR [‡]	61.4	±0.81	61.6	±0.76	83.1	±0.41	83.2	±0.47	27.3	±1.01	30.3	±0.70
	EMAR [†]	62.7	±0.63	62.8	±0.62	82.4	±0.86	84.0	±0.78	45.1	±1.48	46.4	±2.00
	MLLRE	59.8	±0.91	59.8	±0.94	85.2	±0.25	85.5	±0.31	36.4	± 0.66	38.0	± 0.58
	CML (ours)	62.9	± 0.62	63.0	± 0.59	86.5	± 0.22	86.9	± 0.28	43.7	±0.83	45.3	±0.72

Table 1: The average accuracy Acc_a and whole accuracy Acc_w with error bounds by 0.95 confidence, on the test sets of observed tasks at the final time step, where \dagger and \ddagger indicate the result generated from the source code provided by (??)¹ and (??)² respectively.

Train		100	200	all	
Memory		25	50	50	
EA-EMR	Acc _a	70.7	75.5	74.8	73.9
	Acc _w	53.2	57.4	59.8	59.6
MLLRE	Acc _a	68.4	72.1	70.2	51.0
	Acc _w	51.9	57.8	56.8	47.3
EMAR	Acc _a	60.1	66.7	68.6	74.1
	Acc _w	43.7	51.2	53.8	57.7
CML	Acc _a	73.6	76.4	76.0	58.0
	Acc _w	54.7	60.3	60.2	49.1

Table 2: Experimental results on the impact of the amount and ratio of memory on model performance over Continual-FewRel.

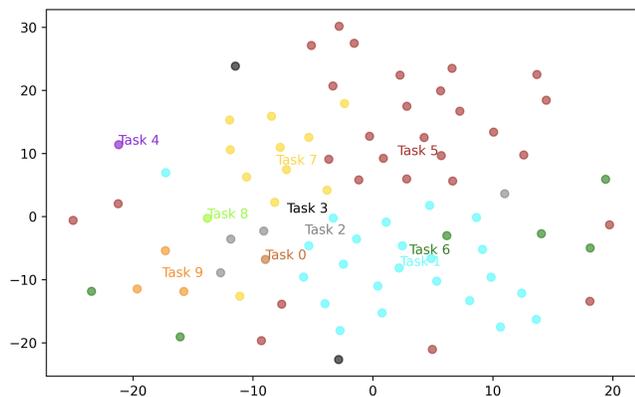


Figure 1: The t-SNE (??) visualization of the encoding of relations generated by the hidden layer in EA-EMR over Continual-FewRel, where each point represents a relationship, each color represents a task, and the position anchored by the label is the center of the task.

ence in their distances from the other relations may explain the difference in their task-specific performance in Table ??, where T_4 does not suffer from catastrophic forgetting but T_8 does. In other words, T_4 is easy whereas T_8 is more difficult.

Similarly, we can observe that T_3 (colored in black) contains only two relations, and it overlaps significantly with the other tasks. Therefore, catastrophic forgetting is more serious on T_3 , i.e., T_3 is difficult. Finally, we can observe that tasks T_1 and T_6 are difficult tasks, as their centroids are very close. They both suffer from serious catastrophic forgetting as can be seen from Table ?. Therefore, in continual learning, the difficulty of a task T may be characterized by the correlation between T and the other observed tasks. In the relation extraction scenario, we define this correlation as the semantic similarity between the relations.

Quantitative Analysis. Based on the above analysis, we hypothesize that a model could achieve better performance if we can measure the similarity of the current task and previous tasks and guide the model to distinguish similar tasks.

For ease of expression, we denote the measure difficulty of each task T_i as the *prior difficulty* D_{prior} , and denote the average forgetting rate F_{avg}^i as the *posterior difficulty* D_{post} . The prior difficulty D_{prior} is the estimated value of knowledge graph embedding, and the posterior difficulty is related to the performance of each model.

Specifically, we use our embedding-based task difficulty function defined in Eq. (??) as the prior difficulty D_{prior} , and average forgetting rate defined in Eq. (??) as the posterior difficulty. Table ?? shows the correlation between the prior difficulty and the posterior difficulty, and thus offers evidence of the effectiveness of our representation learning method.

We can learn from Table ?? four main conclusions: (i) The Pccs of the three models demonstrate that the semantic embedding-based prior difficulty does positive to the for-

task ID	T ₀	T ₁	T ₂	T ₃	T ₄	T ₅	T ₆	T ₇	T ₈	T ₉	Acc _a	Acc _w
run ID	0	1	2	3	4	5	6	7	8	9		
P ₀	88.1	39.0	49.9	59.9	100.0	23.8	28.2	48.5	52.6	53.9	76.8	62.7
P ₁	82.2	43.5	68.8	33.9	100.0	30.7	56.7	39.6	71.1	59.0	73.4	63.9
P ₂	83.0	49.9	66.7	75.9	100.0	26.7	57.4	74.7	58.5	64.1	72.6	61.2
P ₃	91.9	48.3	76.3	77.0	100.0	23.2	54.3	77.0	83.7	50.6	62.8	59.0
P ₄	90.4	44.1	73.6	79.6	100.0	43.5	47.5	69.6	86.7	77.6	77.7	58.4
P ₅	97.0	42.8	71.2	79.9	100.0	46.0	74.9	68.5	80.0	76.8	79.5	61.6
P ₆	98.5	79.2	69.3	75.5	100.0	52.4	82.6	88.1	88.1	80.4	74.6	58.7
P ₇	97.0	80.2	92.1	67.9	100.0	57.1	81.9	87.6	93.3	81.4	67.4	56.3
P ₈	91.9	81.5	91.9	92.7	99.3	68.0	86.1	93.4	92.6	90.8	77.4	56.1
P ₉	100.0	89.7	97.4	96.0	100.0	82.1	91.9	94.7	99.3	93.1	77.7	58.7
μ	92.0	59.8	75.7	73.8	99.9	45.4	66.2	74.2	80.6	72.8	74.0	59.7
σ	6.25	20.06	14.39	17.51	0.22	19.92	20.41	18.50	15.35	14.99	5.26	2.61

Table 3: A case study of EA-EMR on the FewRel dataset, where task ID = $\{T_0, \dots, T_9\}$ represents the fixed ID of each task, run ID = $\{0, \dots, 9\}$ represents the runs with loop offset from 0 to 9, Position = P_0, \dots, P_9 represents the position of a task in a run. μ and δ are the mean and variance of accuracy (read vertically) respectively. Different runs are listed diagonally, for example, when T_0 is at P_1 and run ID is 9, we can obtain the $acc_0^1 = 82.2\%$ at the final time step, and $Acc_a = 76.8\%$, $Acc_w = 62.7\%$. See Appendix C for such result of EMAR, MLLRE and the proposed CML.

		T ₀	T ₁	T ₂	T ₃	T ₄	T ₅	T ₆	T ₇	T ₈	T ₉	PCC _s
	D _{prior}	0.121	0.141	0.168	0.054	0.035	0.186	0.112	0.152	0.146	0.137	-
D _{post}	EA-EMR	0.060	0.098	0.061	0.028	0.001	0.137	0.139	0.060	0.044	0.051	0.559
	MLLRE	0.022	0.078	0.091	0.085	-0.004	0.147	0.060	0.064	0.069	0.075	0.667
	EMAR	0.036	0.016	0.027	0.007	0.006	0.016	0.008	0.026	0.020	0.005	0.499
	CML	-0.002	0.108	0.051	0.065	-0.002	0.113	0.108	0.046	0.070	0.068	0.470

Table 4: Prior difficulty and posterior difficulty of each task, where PCC_s indicates the Pearson correlation coefficient between the estimated prior difficulty by Eq.(?) (i.e., D_{prior}) and the average forgetting rate of each model (i.e., D_{post}).

getting rate indicating which tasks are difficult for a model.

(ii) Comparing the PCC_s values of CML with the other three methods, our proposed CML with the knowledge-based curriculum does alleviate the interference between similar tasks as it achieves the lowest forgetting rate. (iii) For tasks T_0 and T_4 in CML, the average forgetting rate is negative, which means that the accuracy of these tasks decreases as they move towards the end. (iv) Based on the analysis of Table ??, Table ?? (See Appendix C for the other three relevant tables), it is evident that CML improves the overall accuracy and effectively alleviates the order-sensitivity problem. We note that the improvements do come at a cost of moderately decreased accuracy on simple tasks such as T_0 and T_4 .

Conclusion

In this paper, we proposed a novel curriculum-meta learning method to tackle the catastrophic forgetting and order-sensitivity issues in continual relation learning. The construction of the curriculum is based on the notion of task difficulty, which is defined through a novel relation representation learning method that learns from the distribution of domain and range types of relations. Our comprehensive experiments on the three benchmark datasets show that our proposed method outperforms the state-of-the-art models, and is less prone to catastrophic forgetting and less order-sensitive. In future, we will investigate an end-to-end cur-

riculum model and a new dynamic difficulty measurement based on the framework presented in this paper.

Acknowledgments

Research in this paper was partially supported by the National Key Research and Development Program of China under grants (2018YFC0830200, 2017YFB1002801), the Natural Science Foundation of China grants (U1736204), the Judicial Big Data Research Centre, School of Law at Southeast University.