

Adaptive Prompting for Continual Relation Extraction: A Within-Task Variance Perspective

Minh Le^{1*†}, Tien Ngoc Luu^{2†}, An Nguyen The^{3†}, Thanh-Thien Le^{1*}, Trang Nguyen^{1*},
Tung Thanh Nguyen⁴, Linh Ngo Van^{2‡}, Thien Huu Nguyen⁵

¹VinAI Research

²Hanoi University of Science and Technology

³FPT Software AI Center

⁴Moreh Inc.

⁵University of Oregon, Eugene, Oregon, USA

v.minhld12@vinai.io, ngoc.lt204595@sis.hust.edu.vn, annt68@fpt.com, v.thientl3@vinai.io, v.trangnvt2@vinai.io,
tung.nguyen@moreh.com.vn, linhnv@soict.hust.edu.vn, thien@cs.uoregon.edu

Abstract

To address catastrophic forgetting in Continual Relation Extraction (CRE), many current approaches rely on memory buffers to rehearse previously learned knowledge while acquiring new tasks. Recently, prompt-based methods have emerged as potent alternatives to rehearsal-based strategies, demonstrating strong empirical performance. However, upon analyzing existing prompt-based approaches for CRE, we identified several critical limitations, such as inaccurate prompt selection, inadequate mechanisms for mitigating forgetting in shared parameters, and suboptimal handling of cross-task and within-task variances. To overcome these challenges, we draw inspiration from the relationship between prefix-tuning and mixture of experts, proposing a novel approach that employs a prompt pool for each task, capturing variations within each task while enhancing cross-task variances. Furthermore, we incorporate a generative model to consolidate prior knowledge within shared parameters, eliminating the need for explicit data storage. Extensive experiments validate the efficacy of our approach, demonstrating superior performance over state-of-the-art prompt-based and rehearsal-free methods in continual relation extraction.

1 Introduction

Continual Relation Extraction (CRE) involves classifying semantic relationships between entities in text while adapting to an expanding set of relation types. Previous CRE approaches (Zhao et al. 2022; Nguyen et al. 2023; Le et al. 2024d) have successfully addressed the challenge of learning new relations without sacrificing accuracy on previously learned ones by employing memory-based techniques (Shin et al. 2017; Chaudhry et al. 2019). These methods utilize a rehearsal mechanism supported by a memory buffer, enabling the model to revisit and consolidate knowledge

of prior relations while learning new tasks, thereby reducing catastrophic forgetting. Nonetheless, concerns regarding data storage and privacy have prompted the research community to investigate alternative strategies for CRE (Ke and Liu 2022).

To address these limitations, recent advances in Continual Learning (CL) have introduced innovative prompt-based methods (Wang et al. 2022b,a, 2023a; Tran et al. 2024a). Unlike memory-based approaches, these methods eliminate the need for rehearsal, focusing instead on the dynamic insertion of auxiliary parameters, known as prompts, during training. These prompts are adaptable to specific tasks, enabling continual learning without the necessity of data replay. However, our analysis identifies several inherent weaknesses in these prompt-based approaches. Firstly, they lack robust mechanisms to prevent forgetting in shared components, such as the shared Prompt Pool (Wang et al. 2022b), the task-agnostic General Prompt (G-Prompt) (Wang et al. 2022a), or the shared MLP classifier, which can lead to potential performance degradation. Secondly, task-specific prompt-based approaches (Wang et al. 2022a, 2023a) are prone to inaccuracies in prompt selection, leading to a mismatch between training and testing prompts. Lastly, these methods demonstrate limited optimization in managing both cross-task and within-task variance. For instance, Wang et al. (2022b) employ a common prompt pool, where instances from different tasks may frequently share one or more prompts, thereby reducing cross-task variance. This issue becomes particularly prominent in CRE, where instances from different relation classes might frequently have very similar contexts, as shown in the example below:

- "[X] is a professor at [Z university]."
- "[X] is advised by a professor at [Z university]."

In alignment with these approaches, Le et al. (2024a) explores the relationship between Prefix-tuning (Li and Liang 2021), a widely used technique for implementing prompts,

*Work done at Hanoi University of Science and Technology.

†These authors contributed equally.

‡Corresponding Author.

2 Background

2.1 Continual Relation Extraction

Continual Relation Extraction (CRE) is a subfield of continual learning (Hai et al. 2024; Phan et al. 2022; Van et al. 2022) and continual information extraction (Le et al. 2024c; Dao et al. 2024; Tran et al. 2024b). It involves training a model sequentially on a series of tasks $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_T\}$, with each task \mathcal{T}_t associated with a training dataset \mathcal{D}_t and a corresponding set of relations \mathcal{R}_t . Similar to traditional supervised classification frameworks (Ji et al. 2020), each task \mathcal{T}_t consists of \mathcal{N}_t labeled samples, $\mathcal{D}_t = \{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{\mathcal{N}_t}$, where \mathbf{x}_i^t denotes the input data, and y_i^t corresponds to a label from the relation set \mathcal{R}_t . The primary objective in CRE is to train a model that can effectively learn from new tasks while preserving its performance on previously acquired tasks. Upon completing the t -th task, the model should be able to accurately identify the relation of an entity pair from the cumulative relation set $\hat{\mathcal{R}}_t = \bigcup_{i=1}^t \mathcal{R}_i$. Most existing methods in CRE rely on the use of a memory buffer to store samples of previously encountered relations, which raises significant concerns regarding memory and privacy.

2.2 Mixture of Experts

An MoE model consists of N expert networks, denoted by $f_i: \mathbb{R}^D \rightarrow \mathbb{R}^{D_v}$ for $i = 1, \dots, N$, and a gating function, $G: \mathbb{R}^D \rightarrow \mathbb{R}^N$, which dynamically determines the contribution of each expert for a given input \mathbf{x} . The gating function is based on learned score functions, $s_i: \mathbb{R}^D \rightarrow \mathbb{R}$, associated with each expert, resulting in the following formulation:

$$\mathbf{y} = \sum_{j=1}^N G(\mathbf{x})_j \cdot f_j(\mathbf{x}) = \sum_{j=1}^N \frac{\exp(s_j(\mathbf{x}))}{\sum_{\ell=1}^N \exp(s_\ell(\mathbf{x}))} \cdot f_j(\mathbf{x}), \quad (1)$$

where $G(\mathbf{x}) = \text{softmax}(s_1(\mathbf{x}), \dots, s_N(\mathbf{x}))$. Shazeer et al. (2017) introduced Sparse Mixture of Experts (SMoE) architecture as an efficient method to scale up MoE models. This is achieved by utilizing a sparse gating function TopK, which selects only the K experts with the largest affinity scores $s_j(\mathbf{x})$. The TopK function is defined as:

$$\text{TopK}(\mathbf{v}, K)_i = \begin{cases} v_i, & \text{if } v_i \text{ is in the } K \text{ largest elements of } \mathbf{v} \\ -\infty, & \text{otherwise.} \end{cases}$$

Subsequently, the selected experts independently calculate their outputs, and these are linearly combined using their corresponding affinity scores to produce the final prediction:

$$\mathbf{y} = \sum_{j=1}^N \text{softmax}(\text{TopK}(s(\mathbf{x}), K))_j \cdot f_j(\mathbf{x}), \quad (2)$$

where $s(\mathbf{x}) = (s_1(\mathbf{x}), \dots, s_N(\mathbf{x}))$. MoE has gained significant attention for its flexibility and adaptability in fields like large language models (Du et al. 2022; Zhou et al. 2023), computer vision (Riquelme et al. 2021), and multi-task learning (Ma et al. 2018).

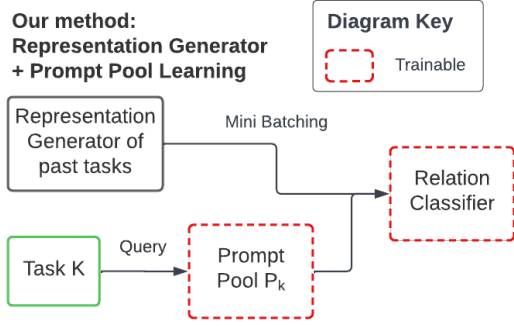


Figure 1: **Overall framework of WAVE-CRE.** To prevent information loss across tasks, we use a task-specific prompt pool \mathbf{P}_t for each task and a representation generator to synthesize past-task information, strengthening the relation classifier’s knowledge retention.

and Mixture of Experts (MoE) models (Jacobs et al. 1991; Jordan and Jacobs 1994). The study demonstrates that self-attention can be seen as embodying multiple MoE models, and that implementing prefix-tuning is analogous to adding new *prefix* experts to these pre-trained MoE models to fine-tune their representations. Building on this foundation, we introduce a novel prompting method, **WAVE-CRE** (Within-Task Variance Awareness for Continual Relation Extraction), designed to address the limitations highlighted earlier, as illustrated in Figure 1. During training, each task is allocated a dedicated Prompt Pool (Wang et al. 2022b), fostering cross-task divergence while capturing intra-task variations. To mitigate catastrophic forgetting during the shared parameter learning process between tasks, we employ a generative model that generates latent data representations for replay. Unlike generating natural language text, learning the underlying distribution and generating continuous latent representations is significantly more feasible. For prompt pool selection, a separate generative model generates uninstructed representations, which are then utilized to train the task predictor. Extensive experiments demonstrate that our method surpasses state-of-the-art prompt-based and rehearsal-free baselines. Therefore, our contributions are as follows:

- We reveal limitations of current prompt-based approaches, including inaccurate prompt selection, inadequate strategies for mitigating forgetting in shared parameters, and suboptimal handling of cross-task and within-task variances.
- To overcome these challenges, we propose a novel prompting method, WAVE-CRE, which leverages task-specific prompt pools and generative models for latent representation.
- Our extensive experimental evaluation demonstrates that WAVE-CRE significantly outperforms state-of-the-art prompt-based and rehearsal-free baselines.

2.3 Prompt-based Methods

Recently, parameter-efficient fine-tuning techniques like Prompt-tuning (Lester, Al-Rfou, and Constant 2021) and Prefix-tuning (Liu et al. 2022; Le et al. 2024b) have become prominent for fine-tuning pre-trained models on downstream tasks. This study focuses on prefix-tuning for prompt implementation, where prompts are passed to several Multi-head Self-attention (MSA) layers in the pre-trained transformer encoder. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D}$ represent the input matrix, where \mathbf{x}_i is the embedding of the i -th token, and D is the embedding dimension. The output of an MSA layer is:

$$\text{MSA}(\mathbf{X}) = \text{Concat}(h_1, \dots, h_m)W_O, \quad (3)$$

$$h_i = \text{Attention}(\mathbf{X}W_i^Q, \mathbf{X}W_i^K, \mathbf{X}W_i^V), \quad (4)$$

for $i = 1, \dots, m$, where W_O , W_i^Q , W_i^K , and W_i^V are projection matrices, and m is the number of heads. In prefix-tuning, a prompt $P \in \mathbb{R}^{L_p \times D}$ of length L_p is divided into $P_k, P_v \in \mathbb{R}^{\frac{L_p}{2} \times D}$. Each head h_i calculation is modified as:

$$\hat{h}_i = \text{Attention}(\mathbf{X}W_i^Q, [P_k; \mathbf{X}]W_i^K, [P_v; \mathbf{X}]W_i^V), \quad (5)$$

where $[\cdot; \cdot]$ denotes the concatenation operation along the sequence length dimension.

Recent research by Le et al. (2024a) has demonstrated that self-attention can be interpreted as a specialized architecture comprising multiple MoE models. The study further suggests that prefix-tuning functions as a mechanism to introduce new experts into these MoE models. Specifically, consider the l -th head in the MSA layer, with output $h_l = [h_{l,1}, \dots, h_{l,N}]^\top \in \mathbb{R}^{N \times D_v}$. Let $\tilde{\mathbf{X}} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_N^\top]^\top \in \mathbb{R}^{N \cdot D}$ represent the concatenation of all input token embeddings. We define N experts $f_j: \mathbb{R}^{N \cdot D} \rightarrow \mathbb{R}^{D_v}$ and N gating functions $G_i: \mathbb{R}^{N \cdot D} \rightarrow \mathbb{R}^N$ with input $\tilde{\mathbf{X}}$ as follows:

$$f_j(\tilde{\mathbf{X}}) = W_l^V{}^\top E_j \tilde{\mathbf{X}} = W_l^V{}^\top \mathbf{x}_j, \quad (6)$$

$$G_i(\tilde{\mathbf{X}}) = \text{softmax}(s_{i,1}(\tilde{\mathbf{X}}), \dots, s_{i,N}(\tilde{\mathbf{X}})),$$

$$s_{i,j}(\tilde{\mathbf{X}}) = \frac{\tilde{\mathbf{X}}^\top E_i^\top W_l^Q W_l^K{}^\top E_j \tilde{\mathbf{X}}}{\sqrt{D_v}}, \quad (7)$$

for $i, j = 1, \dots, N$, where $E_i \in \mathbb{R}^{D \times N \cdot D}$ are matrices such that $E_i \tilde{\mathbf{X}} = \mathbf{x}_i$, and $D_v = \frac{D}{m}$ is the key dimension. Then, from equation (4), the output of the l -th head can be expressed as:

$$\begin{aligned} h_{l,i} &= \sum_{j=1}^N G_i(\tilde{\mathbf{X}})_j \cdot f_j(\tilde{\mathbf{X}}) \\ &= \sum_{j=1}^N \frac{\exp(s_{i,j}(\tilde{\mathbf{X}}))}{\sum_{\ell=1}^N \exp(s_{i,\ell}(\tilde{\mathbf{X}}))} \cdot f_j(\tilde{\mathbf{X}}), \end{aligned} \quad (8)$$

for $i = 1, \dots, N$. From equation (8), we observe that each head h_l in the MSA layer includes N MoE models $h_{l,1}, \dots, h_{l,N}$. This structure is similar to the Multi-gate Mixture of Experts (Ma et al. 2018), where multiple MoE

models leverage the same set of expert networks but employ independent gating functions. Extending this concept, Le et al. (2024a) proposes that prefix-tuning can be viewed as a method for incorporating new experts into these MoE models. New prefix experts and score functions can be defined as:

$$f_{N+j}(\tilde{\mathbf{X}}) = W_l^V{}^\top \mathbf{p}_j^v, \quad (9)$$

$$s_{i,N+j}(\tilde{\mathbf{X}}) = \frac{\tilde{\mathbf{X}}^\top E_i^\top W_l^Q W_l^K{}^\top \mathbf{p}_j^k}{\sqrt{D_v}}, \quad (10)$$

for $i = 1, \dots, N$ and $j = 1, \dots, L$, where $P_k = [\mathbf{p}_1^k, \dots, \mathbf{p}_L^k]^\top$, $P_v = [\mathbf{p}_1^v, \dots, \mathbf{p}_L^v]^\top$, and $L = \frac{L_p}{2}$. From equation (5), the output of the l -th head can be written as $\hat{h}_l = [\hat{h}_{l,1}, \dots, \hat{h}_{l,N}]^\top \in \mathbb{R}^{N \times D_v}$, where

$$\begin{aligned} \hat{h}_{l,i} &= \sum_{j=1}^N \frac{\exp(s_{i,j}(\tilde{\mathbf{X}}))}{\sum_{k=1}^{N+L} \exp(s_{i,k}(\tilde{\mathbf{X}}))} f_j(\tilde{\mathbf{X}}) \\ &+ \sum_{j'=1}^L \frac{\exp(s_{i,N+j'}(\tilde{\mathbf{X}}))}{\sum_{k=1}^{N+L} \exp(s_{i,k}(\tilde{\mathbf{X}}))} f_{N+j'}(\tilde{\mathbf{X}}), \end{aligned} \quad (11)$$

for $i = 1, \dots, N$. These new experts, f_{N+1}, \dots, f_{N+L} , collaborate with the pre-trained experts f_1, \dots, f_N to adapt the model for downstream tasks. Several recent methods have effectively integrated prompting techniques with pre-trained transformer encoders, yielding notable results as demonstrated by L2P (Wang et al. 2022b), DualPrompt (Wang et al. 2022a), and HiDe-Prompt (Wang et al. 2023a).

3 Methodology

Our overall approach, depicted in Figure 1, involves three main stages: (1) Prompt pool learning for a new task; (2) Generative models; and (3) Training the task predictor and relation classifier.

3.1 Task-specific Prompt Pool

In our approach to CRE, we adopt a frozen BERT model (Devlin et al. 2019) as the pre-trained transformer encoder, maintaining consistency with prior studies (Xia et al. 2023; Zhao, Cui, and Hu 2023).

Previous approaches, such as HiDe-Prompt (Wang et al. 2023a), utilize a single prompt per task. This strategy can be compared to utilizing a fixed set of experts for every instance within a given task. However, as detailed in equation (6) and equation (9), the prefix experts encoded in these prompts are considerably simpler, functioning as offset vectors rather than pre-trained experts, which are linear functions of the input. This inherent simplicity implies that a fixed set of prefix experts may lack the necessary flexibility to effectively capture the full range of task variations. To address this limitation, we extend the concept of the prompt pool introduced in L2P (Wang et al. 2022b) by proposing a task-specific prompt pool. For each task t , we introduce a prompt pool \mathbf{P}_t :

$$\mathbf{P}_t = \{(\mathbf{k}_1^{(t)}, P_1^{(t)}), (\mathbf{k}_2^{(t)}, P_2^{(t)}), \dots, (\mathbf{k}_M^{(t)}, P_M^{(t)})\}, \quad (12)$$

where M represents the number of prompts. Each prompt $P_i^{(t)}$ is associated with a learnable key $\mathbf{k}_i^{(t)} \in \mathbb{R}^D$. To

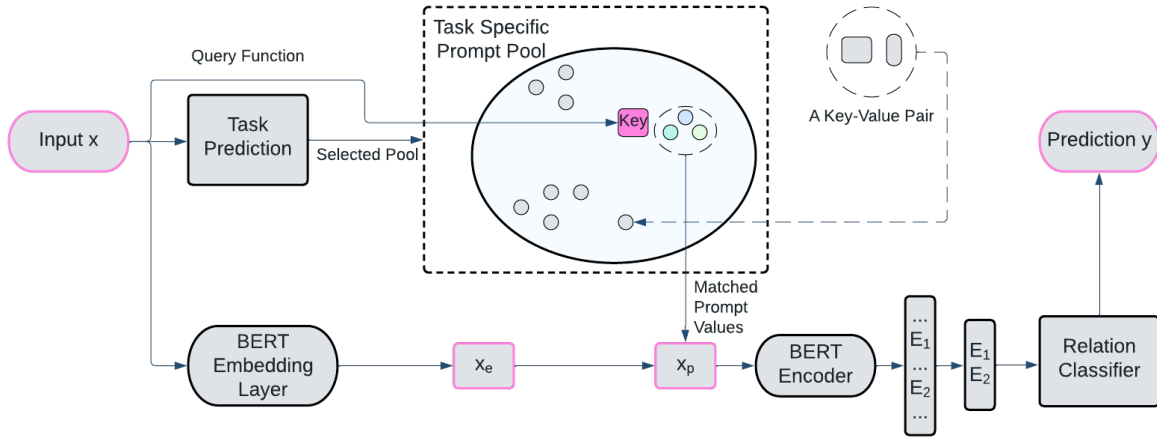


Figure 2: **Data Flow Diagram:** Initially, the task predictor predicts the task identity of the input \mathbf{x} , enabling the selection of the corresponding prompt pool. Subsequently, the input \mathbf{x} queries this prompt pool to identify prompts whose corresponding keys are closest to the *query* $q(\mathbf{x})$. The chosen prompt is then prepended to the embedded input \mathbf{x}_e , creating the prompted input \mathbf{x}_p . The combined \mathbf{x}_p is fed into the BERT Encoder, where the two embeddings corresponding to the positions of the entities E_1 and E_2 are concatenated. Finally, the resulting concatenated embedding is passed to the relation classifier, which predicts the relation label y of the input \mathbf{x} .

facilitate prompt selection, we adopt the same key-query mechanism described in L2P. Specifically, given an input sentence \mathbf{x} , it is first encoded using a pre-trained BERT model to generate a query vector $q(\mathbf{x})$. A scoring function $\gamma : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ then evaluates the match between the query vector and each prompt key (e.g., using cosine similarity). The top K most relevant prompts are selected by optimizing the following objective:

$$\mathbf{K}_{\mathbf{x}} = \underset{S \subseteq \{1, \dots, M\} : |S|=K}{\operatorname{argmin}} \sum_{s \in S} \gamma(q(\mathbf{x}), \mathbf{k}_s^{(t)}), \quad (13)$$

where $\mathbf{K}_{\mathbf{x}}$ denotes a subset of the top- K keys that are specifically chosen for \mathbf{x} .

We strategically set the number of experts within a prompt to one, with $L_p = 2$, resulting in $L = \frac{L_p}{2} = 1$. Using prompts of a larger length would be akin to incorporating more experts per prompt, all of which would share a common prompt key within the prompt pool. However, our approach not only reduces memory costs but also enhances flexibility in selecting experts during training and testing. This configuration allows each expert to adapt to different inputs, providing a more versatile assignment compared to methods that use multiple experts per prompt.

Our proposed architecture enables the assignment of different sets of prompts or experts to specific regions of the input data, guided by the contextual query feature $q(\mathbf{x})$. This design allows each prompt within the pool to selectively focus on the relevant information and patterns necessary for optimal performance in distinct areas of the input domain. As a result, the model is capable of capturing within-task variations effectively. Furthermore, by utilizing a task-specific prompt pool, we reduce the need for parameter sharing, which helps maximize cross-task variance.

Relationship with Sparse Mixture of Experts. Our proposed task-specific prompt pool shares certain similarities with the SMoE architecture in Section 2.2. Specifically, from equation (12), we denote the experts encoded in the prompt pool \mathbf{P}_t as $f_{N+1}^{(t)}, \dots, f_{N+M}^{(t)}$. Unlike the pre-trained experts f_1, \dots, f_N , which are selected by default, we employ sparse selection exclusively for these newly introduced prefix experts. As illustrated in equation (11), each head in the MSA layer when applying prefix-tuning encompasses N MoE models $\hat{h}_{l,1}, \dots, \hat{h}_{l,N}$. The standard approach involves applying the TopK function to each of these N models individually, necessitating the computation of all $N \times M$ score functions $s_{i,N+j}(\tilde{\mathbf{X}})$ for $i = 1, \dots, N$ and $j = 1, \dots, M$. This results in a distinct set of prefix experts selected for each model. Conversely, our strategy leverages the same set of K new experts across all N MoE models using auxiliary score functions defined as:

$$\hat{s}_{i,N+j}(\tilde{\mathbf{X}}) = \gamma(q(\mathbf{x}), \mathbf{k}_j^{(t)}), \quad (14)$$

where $i = 1, \dots, N$ and $j = 1, \dots, M$. This approach only requires the computation of M score functions, as the computation of $\hat{s}_{i,N+j}(\tilde{\mathbf{X}})$ only depends solely on $\mathbf{k}_j^{(t)}$, thereby enabling the efficient and effective selection of K experts from the prompt pool. Although the computation of $q(\mathbf{x})$ might appear to be an added expense, this value is already calculated for the task predictor in Section 3.3 during both the training and testing phases. Therefore, the computation of $q(\mathbf{x})$ can be reused, incurring no additional cost in the prompt selection process.

Optimization Objective. For each new task \mathcal{T}_t , a new prompt pool \mathbf{P}_t is created. During each training step, following the aforementioned strategy, K prompts are selected, and the corresponding prompted embedding feature, de-

noted as \mathbf{x}_p , is inputted to the pre-trained transformer encoder f_r and the final classifier g_ϕ , which is parameterized by ϕ . The objective can be summarized as follows:

$$\min_{\mathbf{P}_t, \phi} \mathcal{L}(g_\phi(f_r(\mathbf{x}_p)), y) + \lambda \sum_{s_i \in \mathbf{K}_x} \gamma(q(\mathbf{x}), \mathbf{k}_{s_i}^{(t)}), \quad (15)$$

where \mathbf{K}_x is obtained with equation (13), γ denotes the cosine similarity function, and λ is a hyperparameter. The first term in equation (15) employs the classification loss, while the second term minimizes the distance between prompt keys and the corresponding query features. Note that only the prompt parameters in \mathbf{P}_t and the final classifier g_ϕ are learned during training task t . The pre-trained BERT model and previous pools $\mathbf{P}_1, \dots, \mathbf{P}_{t-1}$ remain frozen.

3.2 Generative Models for Relation Representation

To effectively retain knowledge acquired from prior tasks, we utilize a generative model that captures the distributions of observed relations, enabling the replay of relation samples. For each relation $r \in \hat{\mathcal{R}}_t$, we maintain a distribution $\mathbf{G}_z^r \sim \mathcal{N}(\boldsymbol{\mu}_z^r, \Sigma_z^r)$. This distribution is obtained by fitting a Gaussian distribution to the set $\mathcal{D}_z^r = \{\mathbf{z}^r = f_r(\mathbf{x}_p^r)\}$, which represents the prompted representation of the input \mathbf{x}^r for relation r :

$$\boldsymbol{\mu}_z^r = \sum_{\mathbf{z}^r} \frac{\mathbf{z}^r}{|\mathcal{D}_z^r|}, \quad \Sigma_z^r = \sum_{\mathbf{z}^r} \frac{(\mathbf{z}^r - \boldsymbol{\mu}_z^r)(\mathbf{z}^r - \boldsymbol{\mu}_z^r)^\top}{|\mathcal{D}_z^r|}. \quad (16)$$

Similarly, the distribution of the query corresponding to each relation denoted as $\mathcal{D}_q^r = \{\mathbf{q}^r = q(\mathbf{x}^r)\}$, is stored as $\mathbf{G}_q^r \sim \mathcal{N}(\boldsymbol{\mu}_q^r, \Sigma_q^r)$:

$$\boldsymbol{\mu}_q^r = \sum_{\mathbf{q}^r} \frac{\mathbf{q}^r}{|\mathcal{D}_q^r|}, \quad \Sigma_q^r = \sum_{\mathbf{q}^r} \frac{(\mathbf{q}^r - \boldsymbol{\mu}_q^r)(\mathbf{q}^r - \boldsymbol{\mu}_q^r)^\top}{|\mathcal{D}_q^r|}. \quad (17)$$

This approach ensures that we have a generative model for each relation, capable of reconstructing the corresponding prompted and query representations for all previously observed relations without storing any instance-specific data. The choice of a Gaussian distribution is motivated by its memory efficiency, as it requires storing only the mean vectors and covariance matrices, thereby minimizing the overall memory footprint. Future works may explore alternative generative models.

3.3 Task Predictor and Relation Classifier

Each task is associated with its own prompt pool, designed to adapt and learn from the samples specific to the task. However, at test time, it becomes crucial to identify which prompt pool corresponds to a new, unseen sample. To tackle this challenge, we introduce a task predictor, denoted as \hat{g}_ψ . This predictor is a feed-forward MLP with an output dimension matching the total number of relations encountered thus far (*i.e.*, $|\hat{\mathcal{R}}_t|$). Once trained, \hat{g}_ψ is capable of predicting the task identity, thereby facilitating the selection of the appropriate prompt pool during testing.

To train the task predictor, we utilize \mathbf{G}_q^r described in Section 3.2 to generate a representation set \mathbf{q}^r for each relation

Algorithm 1: \mathcal{T}_t training process

Input: Training t -th dataset \mathcal{D}_t , current relation set \mathcal{R}_t

Output: Prompt pool \mathbf{P}_t , task predictor ψ , and relation classifier ϕ

```

1: Randomly initialize  $\mathbf{P}_t$ 
2: for  $e_{id} \leftarrow 1$  to training_epoch do
3:   for batch  $\mathbf{x}_B \in \mathcal{D}_t$  do
4:     Update  $\mathbf{P}_t$  and  $g_\phi$  on  $\mathbf{x}_B$  via equation (15)
5:   end for
6: end for
7: Update  $\hat{\mathcal{R}}_t \leftarrow \hat{\mathcal{R}}_{t-1} \cup \mathcal{R}_t$ 
8: for each  $r \in \mathcal{R}_t$  do
9:    $\mathcal{D}_q^r \leftarrow \emptyset, \mathcal{D}_z^r \leftarrow \emptyset$ 
10:  for batch  $\mathbf{x}_B \in \mathcal{D}_t^r$  do
11:    Update  $\mathcal{D}_q^r, \mathcal{D}_z^r$ 
12:  end for
13:  Fit  $\mathbf{G}_z^r$  to  $\mathcal{D}_z^r$  via equation (16)
14:  Fit  $\mathbf{G}_q^r$  to  $\mathcal{D}_q^r$  via equation (17)
15: end for
16: Train the task predictor  $\psi$  via equation (18)
17: Train the relation classifier  $\phi$  via equation (19)
18: return  $\mathbf{P}_t, \phi, \psi$ 

```

$r \in \hat{\mathcal{R}}_t$. The task predictor is trained with the cross-entropy loss function defined as:

$$\mathcal{L}(\psi) = \sum_{r \in \hat{\mathcal{R}}_t} \sum_{\mathbf{q} \sim \mathbf{G}_q^r} -\log \frac{\exp(\hat{g}_\psi(\mathbf{q})[r])}{\sum_{r' \in \hat{\mathcal{R}}_t} \exp(\hat{g}_\psi(\mathbf{q})[r'])}. \quad (18)$$

While our approach shares similarities with HiDe-Prompt (Wang et al. 2023a) in utilizing an additional MLP head, a key distinction lies in how relations are treated. HiDe-Prompt categorizes all relations within a task as a single class in the cross-entropy loss. This strategy can be suboptimal, as the resulting classes may lack semantic significance and their meaning can depend on the sequence in which tasks are presented during training.

In a manner analogous to the training of the task predictor, we train the relation classifier g_ϕ using \mathbf{G}_z^r with the same cross-entropy loss function defined as follows:

$$\mathcal{L}(\phi) = \sum_{r \in \hat{\mathcal{R}}_t} \sum_{\mathbf{z} \sim \mathbf{G}_z^r} -\log \frac{\exp(g_\phi(\mathbf{z})[r])}{\sum_{r' \in \hat{\mathcal{R}}_t} \exp(g_\phi(\mathbf{z})[r'])}. \quad (19)$$

This approach helps to mitigate catastrophic forgetting in the shared classification head without requiring the storage of samples from previous relations. For a detailed overview of the training process, please refer to Algorithm 1. The data flow diagram illustrating the inference process is provided in Figure 2.

4 Experiments

4.1 Experimental Settings

Datasets. To evaluate the effectiveness of WAVE-CRE and the baseline models, we utilize two popular datasets:

- **FewRel** (Han et al. 2018) contains 80 relation types with a total of 56,000 samples. Following the configurations

FewRel										
Model	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
EA-EMR	89.0	69.0	59.1	54.2	47.8	46.1	43.1	40.7	38.6	35.2
RP-CRE	97.9	92.7	91.6	89.2	88.4	86.8	85.1	84.1	82.2	81.5
CRL	98.2	94.6	92.5	90.5	89.4	87.9	86.9	85.6	84.5	83.1
CRE-DAS	98.1	<u>95.8</u>	<u>93.6</u>	91.9	<u>91.1</u>	89.4	88.1	86.9	85.6	84.2
CDec+ACA	98.4	95.4	93.2	<u>92.1</u>	91.0	<u>89.7</u>	<u>88.3</u>	<u>87.4</u>	<u>86.4</u>	<u>84.8</u>
L2P	97.4	90.8	83.6	76.5	68.9	64.1	61.0	57.4	50.1	44.6
EPI	98.3	89.9	84.0	79.9	76.5	73.1	70.1	67.0	64.5	61.8
HiDe-Prompt	95.5	89.4	86.0	85.7	87.8	84.2	75.9	75.1	70.3	67.2
WAVE-CRE	97.9	95.5	93.6	92.4	91.1	90.2	88.7	87.6	86.5	85.0

TACRED										
Model	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
EA-EMR	47.5	40.1	38.3	29.9	24	27.3	26.9	25.8	22.9	19.8
RP-CRE	97.6	90.6	86.1	82.4	79.8	77.2	75.1	73.7	72.4	72.4
CRL	<u>97.7</u>	93.2	89.8	84.7	84.1	81.3	80.2	79.1	79.0	78.0
CRE-DAS	<u>97.7</u>	<u>94.3</u>	<u>92.3</u>	<u>88.4</u>	<u>86.6</u>	<u>84.5</u>	<u>82.2</u>	<u>81.1</u>	<u>80.1</u>	<u>79.1</u>
CDec+ACA	<u>97.7</u>	92.8	91.0	86.7	85.2	82.9	80.8	80.2	78.8	78.6
L2P	96.9	88.2	73.8	68.6	66.3	63.1	60.4	59.1	56.8	54.8
EPI	97.5	90.7	82.7	76.7	74.0	72.3	68.2	66.5	65.1	63.4
HiDe-Prompt	97.3	92.8	86.2	82.6	80.6	80.4	75.8	73.7	72.9	72.6
WAVE-CRE	98.4	94.3	91.6	87.8	85.7	83.5	81.3	80.4	79.5	78.7

Table 1: Average accuracy (%) of all methods across learning stages for FewRel and TACRED dataset. The best accuracy scores under the rehearsal-free and rehearsal-based setting are in **bold** and underlined, respectively.

Task-Incremental Learning - TACRED										
Model	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
WAVE-CRE	98.4	95.5	94.2	94.1	92.7	89.9	88.3	87.6	86.5	85.2
w/o Prompt Pool	96.8	94.0	92.9	91.5	90.6	88.0	86.1	84.9	84.4	83.4

Table 2: Detailed analysis of WAVE-CRE with task-specific prompt pool in the task-incremental learning scenario of TACRED. We report the average accuracy across different stages. The best accuracy scores are in **bold**.

outlined in Wang et al. (2019), we split it into 10 non-overlapping sub-datasets.

- **TACRED** (Zhang et al. 2017) consists of 42 relations and 106,264 samples. We adopt the experimental settings proposed by Cui et al. (2021) to partition the dataset into 10 distinct sub-datasets.

Baselines. We compare WAVE-CRE with recent rehearsal-free and prompt-based continual learning methods including L2P (Wang et al. 2022b), HiDe-Prompt (Wang et al. 2023a), and EPI (Wang et al. 2023b). As these methods were originally designed for computer vision, we re-implemented them for CRE using BERT (Devlin et al. 2019) as the encoder. Additionally, we compare our method with rehearsal-based CRE baselines including EA-EMR (Wang et al. 2019), RP-CRE (Cui et al. 2021), CRL (Zhao et al. 2022), CRE-DAS (Zhao, Cui, and Hu 2023), and CDec+ACA (Xia et al. 2023).

Implementation Details. In this work, we used a single NVIDIA A100 for all methods. We tune the hyperparameters for the proposed model using random search. We maintained a consistent size for the prompt pool M across all tasks. For baselines, we follow the identical experimental settings employed by Zhao et al. (2022) to ensure fair comparisons. Our proposed model has in total 114M parameters. Since we froze the BERT model, the number of learnable pa-

rameters is thus only 3.8M. Training on the FewRel dataset took approximately 7 hours, while for the TACRED dataset, it took approximately 3 hours using our method.

Evaluation Metrics. We use the same performance measures (mean accuracy on 5 different random seeds) as in prior work (Zhao et al. 2022) for fair comparison.

4.2 Main Results

Table 1 summarizes the performance of all methods on FewRel and TACRED datasets. We begin by comparing WAVE-CRE with rehearsal-free and prompt-based methods for CRE. Notably, among all rehearsal-free methods, WAVE-CRE consistently outperforms across different stages of training on both datasets. Particularly on the last task \mathcal{T}_{10} , L2P and EPI exhibit substantially lower performance, with a gap of up to 15% in final average accuracy compared to our method. This substantial difference underscores the limitations of these existing approaches in addressing catastrophic forgetting across diverse domains. While HiDe-Prompt shows some improvements, it still experiences performance losses of over 15% and 6% on FewRel and TACRED, respectively. These losses can be attributed to insufficient task-identity inference techniques, as discussed in Section 4.3.

Furthermore, we evaluate WAVE-CRE against recent suc-

Task Incremental Learning - TACRED											
L	K	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
8	1	97.2	94.8	93.1	92.3	90.5	87.8	84.6	84.1	83.4	84.2
4	2	97.7	95.5	94.9	91.2	90.7	88.0	86.5	86.3	85.3	84.1
2	4	96.5	95.1	93.2	92.1	91.2	88.9	85.3	85.2	84.6	84.0
1	8	98.4	95.5	94.2	94.1	92.7	89.9	88.3	87.6	86.5	85.2

Table 3: Detailed analysis of the impact of the number of experts within a prompt. We report the average accuracy across different stages on TACRED in the task incremental learning scenario. The best accuracy scores are in **bold**.

FewRel										
Model	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
EPI	64.5	62.8	64.4	64.9	64.5	64.4	64.4	59.1	61.1	56.6
HiDe-Prompt	76.7	81.7	80.6	80.6	80.2	78.9	77.8	83.1	80.3	81.0
WAVE-CRE	88.5	86.2	86.9	86.4	85.2	86.9	87.7	86.0	85.4	82.5
TACRED										
Model	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5	\mathcal{T}_6	\mathcal{T}_7	\mathcal{T}_8	\mathcal{T}_9	\mathcal{T}_{10}
EPI	64.2	58.8	66.2	55.3	64.0	62.1	67.8	62.0	62.4	62.5
HiDe-Prompt	70.2	72.0	73.8	75.9	68.6	79.8	75.2	71.2	68.5	64.9
WAVE-CRE	83.8	79.5	84.2	76.2	80.4	74.8	79.7	69.3	83.7	81.5

Table 4: Task prediction precision for each task (%) at testing time after training the 10-th task, in comparison with some rehearsal-free methods. The best accuracy scores are in **bold**.

successful CRE methods, all of which are rehearsal-based. Remarkably, without retaining training data or directly fine-tuning BERT, WAVE-CRE achieves results nearly equivalent to these state-of-the-art baselines on both datasets. Notably, in the FewRel dataset, our method surpasses the latest rehearsal-based methods on the last task. This highlights the significance and effectiveness of our approach.

4.3 Detailed Analysis

Task-specific Prompt Pool. To illustrate the efficacy of task-specific prompt pools in capturing within-task variations compared to a single prompt approach, we conducted experiments in the task incremental learning setting (van de Ven, Tuytelaars, and Tolias 2022). In this setting, task identities are provided, ensuring the prompt pool remains identical during training and testing. Table 2 compares models trained with and without the task-specific prompt pool. Here, “w/o Prompt Pool” represents using only a single task-specific prompt per task. Our method, WAVE-CRE, trained with the prompt pool, shows a 1.8% improvement on the final task, demonstrating its effectiveness in enhancing the model’s ability to capture within-task variations.

Number of Experts per Prompt L . Similarly, we investigate the effect of different values of L in the task incremental learning scenario. The number of prompts selected, K , was chosen to ensure a fair comparison by keeping the total number of experts across experiments equal. The results are presented in Table 3. As discussed in Section 3.1, setting $L = 1$ allows for flexibility in expert selection and increases the model’s expressiveness, as each expert has its own key. This is empirically shown, as our proposal achieves the best performance among different values of L .

Detailed Analysis of Task Predictor. We evaluate our task-ID prediction technique against the baselines EPI and

HiDe-Prompt, with results summarized in Table 4. EPI uses BERT for task identification and Mahalanobis distance to select task-specific parameters, assuming pre-trained representations are well-separated—an unrealistic assumption given their generic origin, leading to poor prediction accuracy. In contrast, our method trains a dedicated task predictor on synthesized query representations, significantly improving accuracy. HiDe-Prompt adopts a similar strategy and improves over EPI but groups all relations within a task as a single class, which can be suboptimal (Section 3.3). WAVE-CRE addresses this by using a task predictor with output dimensions corresponding to the number of relations, achieving up to 10% improvements on both datasets, showcasing our strategy’s effectiveness.

5 Conclusion

In this work, we propose a novel framework called WAVE-CRE for rehearsal-free continual relation extraction. Our contributions focus on generating representations for replay, precise task prediction, and optimizing within-task and cross-task variability via prompting. These strategies address limitations of current state-of-the-art prompt-based baselines for continual learning. Through extensive benchmarks, we show our model consistently outperforms existing rehearsal-free methods and achieves competitive results with advanced CRE methods. While our methods mitigate catastrophic forgetting, challenges remain. Retaining knowledge of past tasks is difficult, as seen in prior works. Despite improvements using prompt pools for task-specific knowledge, forgetfulness occurs when pools are not properly utilized during testing. Additionally, while prompt pools enhance expressiveness, the current prefix-tuning experts are relatively simple. Future work could explore more complex expert designs to improve the model.

References

- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2019. Efficient Lifelong Learning with A-GEM. In *International Conference on Learning Representations (ICLR)*.
- 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*, 232–243. Online.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4171–4186.
- Du, N.; Huang, Y.; Dai, A. M.; Tong, S.; Lepikhin, D.; Xu, Y.; Krikun, M.; Zhou, Y.; Yu, A. W.; Firat, O.; Zoph, B.; Fedus, L.; Bosma, M. P.; Zhou, Z.; Wang, T.; Wang, Y. E.; Webster, K.; Pellat, M.; Robinson, K.; Meier-Hellstern, K. S.; Duke, T.; Dixon, L.; Zhang, K.; Le, Q. V.; Wu, Y.; Chen, Z.; and Cui, C. 2022. GLaM: Efficient Scaling of Language Models with Mixture-of-Experts. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 5547–5569. PMLR.
- 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.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive mixtures of local experts. *Neural Computation*, 3.
- Ji, B.; Yu, J.; Li, S.; Ma, J.; Wu, Q.; Tan, Y.; and Liu, H. 2020. Span-based Joint Entity and Relation Extraction with Attention-based Span-specific and Contextual Semantic Representations. In Scott, D.; Bel, N.; and Zong, C., eds., *Proceedings of the 28th International Conference on Computational Linguistics*, 88–99. Barcelona, Spain (Online): International Committee on Computational Linguistics.
- Jordan, M. I.; and Jacobs, R. A. 1994. Hierarchical mixtures of experts and the EM algorithm. *Neural computation*, 6(2): 181–214.
- Ke, Z.; and Liu, B. 2022. Continual learning of natural language processing tasks: A survey. *arXiv preprint arXiv:2211.12701*.
- 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, M.; Nguyen, C.; Nguyen, H.; Tran, Q.; Le, T.; and Ho, N. 2024b. Revisiting Prefix-tuning: Statistical Benefits of Reparameterization among Prompts. *arXiv preprint arXiv:2410.02200*.
- Le, T.-T.; Dao, V.; Nguyen, L.; Nguyen, T.-N.; Ngo, L.; and Nguyen, T. 2024c. 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. 2024d. Continual relation extraction via sequential multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 18444–18452.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045–3059. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. *arXiv:2101.00190*.
- Liu, X.; Ji, K.; Fu, Y.; Tam, W.; Du, Z.; Yang, Z.; and Tang, J. 2022. P-Tuning: Prompt Tuning Can Be Comparable to Fine-tuning Across Scales and Tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 61–68. Dublin, Ireland: Association for Computational Linguistics.
- Ma, J.; Zhao, Z.; Yi, X.; Chen, J.; Hong, L.; and Chi, E. H. 2018. Modeling Task Relationships in Multi-task Learning with Multi-gate Mixture-of-Experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18, 1930–1939*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355520.
- 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.
- Riquelme, C.; Puigcerver, J.; Mustafa, B.; Neumann, M.; Jenatton, R.; Pinto, A. S.; Keyesers, D.; and Houlsby, N. 2021.

- Scaling Vision with Sparse Mixture of Experts. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 8583–8595.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *International Conference on Learning Representations (ICLR)*.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Tran, Q.; 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 preprint arXiv:2410.04327*.
- Tran, Q.; Thanh, N.; Anh, N.; Hai, N.; Le, T.; Ngo, L.; and Nguyen, T. 2024b. 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 de Ven, G. M.; Tuytelaars, T.; and Tolias, A. S. 2022. Three types of incremental learning. *Nature Machine Intelligence*, 4: 1185–1197.
- Wang, H.; Xiong, W.; Yu, M.; Guo, X.; Chang, S.; and Wang, W. Y. 2019. Sentence Embedding Alignment for Lifelong Relation Extraction. 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)*, 796–806. Minneapolis, Minnesota: Association for Computational Linguistics.
- Wang, L.; Xie, J.; Zhang, X.; Huang, M.; Su, H.; and Zhu, J. 2023a. Hierarchical Decomposition of Prompt-Based Continual Learning: Rethinking Obscured Sub-optimality. *Advances in Neural Information Processing Systems*.
- Wang, Z.; Liu, Y.; Ji, T.; Wang, X.; Wu, Y.; Jiang, C.; Chao, Y.; Han, Z.; Wang, L.; Shao, X.; and Zeng, W. 2023b. Rehearsal-free Continual Language Learning via Efficient Parameter Isolation. 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*, 10933–10946. Association for Computational Linguistics.
- Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J.; et al. 2022a. Dual-Prompt: Complementary Prompting for Rehearsal-free Continual Learning. *European Conference on Computer Vision*.
- Wang, Z.; Zhang, Z.; Lee, C.-Y.; Zhang, H.; Sun, R.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022b. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 139–149.
- Xia, H.; Wang, P.; Liu, T.; Lin, B.; Cao, Y.; and Sui, Z. 2023. Enhancing Continual Relation Extraction via Classifier Decomposition. In *Proceedings of the 61th Annual Meeting of the 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 *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, 35–45.
- 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.
- Zhao, W.; Cui, Y.; and Hu, W. 2023. Improving Continual Relation Extraction by Distinguishing Analogous Semantics. In *Proceedings of the 61th Annual Meeting of the Association for Computational Linguistics*.
- Zhou, Y.; Du, N.; Huang, Y.; Peng, D.; Lan, C.; Huang, D.; Shakeri, S.; So, D. R.; Dai, A. M.; Lu, Y.; Chen, Z.; Le, Q. V.; Cui, C.; Laudon, J.; and Dean, J. 2023. Brainformers: Trading Simplicity for Efficiency. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 42531–42542. PMLR.