

Continual Learning for Text Classification with Information Disentanglement Based Regularization

Yufan Huang*, Yanzhe Zhang*, Jiaao Chen, Xuezhi Wang¹, Diyi Yang

Georgia Institute of Technology, ¹Google

{yhuang704, jiaaochen, dyang888}@gatech.edu, ¹xuezhiw@google.com

Abstract

Continual learning has become increasingly important as it enables NLP models to constantly learn and gain knowledge over time. Previous continual learning methods are mainly designed to preserve knowledge from previous tasks, without much emphasis on how to well generalize models to new tasks. In this work, we propose an *information disentanglement* based regularization method for continual learning on text classification. Our proposed method first disentangles text hidden spaces into representations that are generic to all tasks and representations specific to each individual task, and further regularizes these representations differently to better constrain the knowledge required to generalize. We also introduce two simple auxiliary tasks: next sentence prediction and task-id prediction, for learning better generic and specific representation spaces. Experiments conducted on large-scale benchmarks demonstrate the effectiveness of our method in continual text classification tasks with various sequences and lengths over state-of-the-art baselines.

1 Introduction

Computational systems in real world scenarios face changing environment frequently, and thus are often required to learn continually from dynamic streams of data building on what was learnt before (Parisi et al., 2019). While being an intrinsic nature of human to continually acquire and transfer knowledge throughout lifespans, most machine learning models often suffer from *catastrophic forgetting*: when learning on new tasks, models dramatically and rapidly forget knowledge from previous tasks (McCloskey and Cohen, 1989). As a result, Continual Learning (CL) (Ring, 1998; Thrun, 1998) has received more attention recently as it can enable models to perform positive transfer (Perkins et al., 1992) as well as remember previously seen tasks.

A growing body of research has been conducted to equip neural networks with the ability of continual learning abilities (Kirkpatrick et al., 2017; Lopez-Paz and Ranzato, 2017; Aljundi et al., 2018). Existing continual learning methods on NLP tasks can be broadly categorized into two classes: purely *replay* based methods (d’Autume et al., 2019; Sun et al., 2019; Holla et al., 2020) where examples from previous tasks are stored and re-trained during the learning of the new task to retain old information, and *regularization* based methods (Wang et al., 2019; Han et al., 2020) where constraints are added on model parameters to prevent them from changing too much while learning new tasks. The former usually stores an extensive amount of data from old tasks (d’Autume et al., 2019) or trains language models based on task identifiers to generate sufficient examples (Sun et al., 2019), which significantly increases memory costs and training time. While the latter utilizes previous examples efficiently via the constraints added on text hidden space or model parameters, it generally views them as equally important and regularize them to the same extent (Wang et al., 2019), making it hard for models to differentiate informative representation that needs to be retained from ones that need a large degree of updates. However, we argue that when learning new tasks, task generic information and task specific information should be treated differently, as these generic representation might function consistently while task specific representations might need to be changed significantly.

To this end, we propose an information disentanglement based regularization method for continual learning on text classification. Specifically, we first disentangle the text hidden representation space into a **task generic** space and a **task specific** space using two auxiliary tasks: next sentence prediction for learning task generic information and task identifier prediction for learning task specific representations. When training on new tasks, we

*Equal contribution.

constrain the task generic representation to be relatively stable and representations of task specific aspects to be more flexible. To further alleviate catastrophic forgetting without much increases of memory and training time, we propose to augment our regularization-based methods by storing and replaying only a small amount of representative examples (e.g., 1% samples selected by memory selection rules like K-Means (MacQueen et al., 1967)). To sum up, our contributions are threefold:

- We propose an information disentanglement based regularization method for continual text classification, to better learn and constrain task generic and task specific knowledge.
- We augment the regularization approach with a memory selection rule that requires only a small amount of replaying examples.
- Extensive experiments conducted on five benchmark datasets demonstrate the effectiveness of our proposed methods compared to state-of-the-art baselines.

2 Related work

Continual Learning Existing continual learning research can be broadly divided into four categories: (i) replay-based method, which remind models of information from seen tasks via experience replay (d’Autume et al., 2019), distillation (Rebuffi et al., 2017), representation alignment (Wang et al., 2019) or optimization constraints (Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018b) using examples sampled from previous tasks (Rebuffi et al., 2017; d’Autume et al., 2019) or synthesized with generative models (Shin et al., 2017; Sun et al., 2019); (ii) regularization-based method, which constrains model’s output (Li and Hoiem, 2017), hidden space (Rannen et al., 2017), or parameters (Lopez-Paz and Ranzato, 2017; Zenke et al., 2017; Aljundi et al., 2018) from changing too much to retain learned knowledge; (iii) architecture-based method, where different tasks are associated with different components of the overall model to directly minimize the interference between new tasks and old tasks (Rusu et al., 2016; Mallya and Lazebnik, 2018); (iv) meta-learning-based method, which learns robust model initialization (Obamuyide and Vlachos, 2019) or data representations (Javed and White, 2019; Holla et al., 2020; Wang et al., 2020) before training on task sequences to alleviate forgetting.

Among these different approaches, replay-based methods and regularization-based methods have been widely applied to NLP tasks to enable large pre-trained models (Devlin et al., 2018; Radford et al., 2019) to continually acquire novel world knowledge from streams of textual data without forgetting the already learned knowledge. For instance, replaying examples have shown promising performance for text classification (d’Autume et al., 2019; Sun et al., 2019; Holla et al., 2020), relation extraction (Wang et al., 2019; Obamuyide and Vlachos, 2019; Han et al., 2020) and question answering (d’Autume et al., 2019; Sun et al., 2019; Wang et al., 2020). However, they often suffer from large memory costs or considerable training time, due to the requirements of storing an extensive amount of texts (d’Autume et al., 2019; Wang et al., 2019) or training language models to generate a sufficient number of examples (Sun et al., 2019). Recently, regularization-based methods (Wang et al., 2019) have also been applied to directly constrain knowledge deposited in model parameters without abundant rehearsal examples. Despite better efficiency compared to replay-based methods, current regularization-based approaches often fail to generalize well to new tasks as they treat and constrain all the information *equally* and thus limit the needed updates for parameters that are specific to different tasks. To overcome these limitations, we propose to first distinguish hidden spaces that need to be retained from those that need to be updated substantially through information disentanglement, and then regularize different spaces *separately*, to better remember previous knowledge as well as transfer to new tasks. In addition, we enhance our regularization method by replaying only a limited amount of examples selected by K-means as the memory selection rule.

Textual Information Disentanglement Our work is related to information disentanglement for text data, which has been extensively explored in generation tasks like style transfer (Fu et al., 2017; Zhao et al., 2018; Romanov et al., 2018; Li et al., 2020), where text hidden representations are often disentangled into sentiment (Fu et al., 2017; John et al., 2018), content (Romanov et al., 2018; Bao et al., 2019) and syntax (Bao et al., 2019) information through supervised learning from pre-defined labels (John et al., 2018) or unsupervised learning with adversarial training (Fu et al., 2017; Li et al., 2020). Building on these

prior works, we differentiate task generic space from task specific space via supervision from two simple yet effective auxiliary tasks: next sentence prediction and task identifier prediction.

3 Problem Formulation

In this work, we focus on continual learning for a sequence of text classification tasks $\{T_1, \dots, T_n\}$, where we learn a model $f_\theta(\cdot)$, θ is a set of parameters shared by all tasks and each task T_i contains a different set of sentence-label training pairs, $(x_{1:m}^i, y_{1:m}^i)$. After learning all tasks in the sequence, we seek to minimize the generalization error on all tasks (Parisi et al., 2019):

$$R(f_\theta) = \sum_{i=1}^n \mathbb{E}_{(x^i, y^i) \sim T_i} \mathcal{L}(f_\theta(x^i), y^i)$$

We use two commonly-used techniques for this problem setting in our proposed model:

- **Regularization:** in order to preserve knowledge stored in the model, regularization is a constraint added to model output (Li and Hoiem, 2017), hidden space (Zenke et al., 2017) and parameters (Lopez-Paz and Ranzato, 2017; Zenke et al., 2017; Aljundi et al., 2018) to prevent them from changing too much while learning new tasks.
- **Replay:** when learning new tasks, Experience Replay (Rebuffi et al., 2017) is commonly used to recover knowledge from previous tasks, where a memory buffer is first adopted to store seen examples from previous tasks and then the stored data is replayed with the training set for the current task. Formally, after training on task $t-1$ ($t \geq 2$), $\gamma|S_{t-1}|$ examples are *randomly sampled* from the t -th training set S_{t-1} into the memory buffer \mathcal{M} , where $0 \leq \gamma \leq 1$ is the store ratio. Data from \mathcal{M} is then merged with the t -th training set S_t when learning from task t .

4 Method

In continual learning, the model needs to adapt to new tasks quickly while maintaining the ability to recover information from previous tasks, hence not all information stored in the hidden representation space should be treated equally. For example, syntactic knowledge might be shared globally across all tasks, like the ability to recognize word order

and grammar; while some knowledge is unique to certain tasks, and should be preserved separately. This key observation motivates us to propose an *information-disentanglement* based regularization for continual text classification to retain shared knowledge while adapting specific knowledge to streams of tasks (Section 4.1). We also incorporate a small set of representative replay samples to alleviate catastrophic forgetting (Section 4.3). Our model architecture is shown in Figure 1.

4.1 Information Disentanglement (ID)

This section describes how to disentangle sentence representations into task generic space and task specific space, and how separate regularizations are imposed on them for continual text classification. Formally, for a given sentence x , we first use a multi-layer encoder $B(\cdot)$, e.g., BERT (Devlin et al., 2018), to get the hidden representations r which contain both task generic and task specific information. Then we introduce two disentanglement networks $G(\cdot)$ and $S(\cdot)$ to extract the generic representation g and specific representation s from r . For new tasks, we learn the classifiers by utilizing information from both spaces, and we allow different spaces to change to different extents to best retain knowledge from previous tasks.

Task Generic Space Task generic space is the hidden space containing information generic to different tasks in a task sequence. During switching from one task to another, the generic information should roughly remain the same, e.g., syntactic knowledge should not change too much across the learning process of a sequence of tasks. To extract task generic information g from hidden representations r , we leverage the *next sentence prediction* task (Devlin et al., 2018) to learn the generic information extractor $G(\cdot)$. More specifically, we insert a [SEP] token into each training example during tokenization to form a sequence pair labeled *IsNext*, and switch the first sequence and the second sequence to form a sentence pair labeled *NotNext*. In order to distinguish *IsNext* pairs and *NotNext* pairs, extractor $G(\cdot)$ needs to learn the context dependencies between two segments, which is beneficial to understand every example and generic to any individual task.

Denote \tilde{x} as the *NotNext* example corresponding to x (*IsNext*), and $l \in \{0, 1\}$ as the label for next sentence prediction. We build a sentence relation

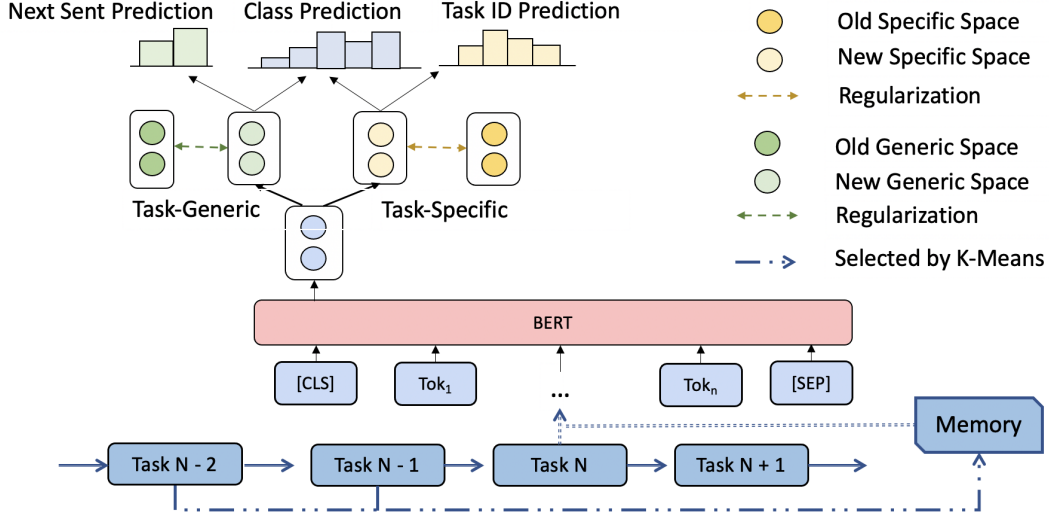


Figure 1: Our proposed model architecture. We disentangle the hidden representation into a task generic space and a task specific space via different induction biases. When training on new tasks, different spaces are regularized separately. Also, a small portion of previous data is stored and replayed.

predictor $f_{n.sp}$ on the generic feature extractor $G(\cdot)$:

$$\mathcal{L}_{n.sp} = \mathbb{E}_{x \in S_t \cup \mathcal{M}} (\mathcal{L}(f_{n.sp}(G(B(x))), 0) + \mathcal{L}(f_{n.sp}(G(B(\tilde{x}))), 1))$$

where \mathcal{L} is the cross entropy loss, \mathcal{M} is the memory buffer and S_t is the t -th training set.

Task Specific Space Models also need task specific information to perform well over each task. For example, on sentiment classification words like “good” or “bad” could be very informative, but they might not generalize well for tasks like topic classification. Thus we employ a simple task-identifier prediction task on the task specific representation s , which means for any given example we want to distinguish which task this example belongs to. This simple auxiliary setup will encourage s to embed different information from different tasks. The loss for task-identifier predictor f_{task} is:

$$\mathcal{L}_{task} = \mathbb{E}_{(x,z) \in S_t \cup \mathcal{M}} \mathcal{L}(f_{task}(S(B(x))), z)$$

where z is the corresponding task id for x .

Text Classification To adapt to the t -th task, we combine the task generic representation $g = G(B(x))$ and task specific representation $s = S(B(x))$ to perform text classification, where we minimize the cross entropy loss:

$$\mathcal{L}_{cls} = \mathbb{E}_{(x,y) \in S_t \cup \mathcal{M}} \mathcal{L}(f_{cls}(g \circ s), y)$$

Here y is the corresponding class label for x , $f_{cls}(\cdot)$ is the class predictor. \circ denotes the concatenation of the two representations.

4.2 ID Based Regularization

To further prevent severe distortion when training on new tasks, we employ regularization on both generic representations g and specific representations s . Different from previous approaches (Li and Hoiem, 2017; Zenke et al., 2017; Aljundi et al., 2018) which treat all the spaces equally, we allow regularization to different extents on g and s as knowledge in different spaces should be preserved separately to encourage both more positive transfer and less forgetting. Specifically, before training all the modules on task t , we first compute the generic representations and specific representations of all sentences x from the training set S_t of current task t and memory buffer \mathcal{M}_t , using the trained $B^{t-1}(\cdot)$, $G^{t-1}(\cdot)$ and $S^{t-1}(\cdot)$ from previous task $t-1$: $g^{t-1} = G^{t-1}(B^{t-1}(x))$, and $s^{t-1} = S^{t-1}(B^{t-1}(x))$ to hoard the knowledge from previous models. The computed generic and specific representations are saved. During the learning from training pairs from task t , we impose two regularization losses separately:

$$\mathcal{L}_{reg}^g = \mathbb{E}_{x \in S_t \cup \mathcal{M}_t} \|G^{t-1}(B^{t-1}(x)) - G(B(x))\|_2$$

$$\mathcal{L}_{reg}^s = \mathbb{E}_{x \in S_t \cup \mathcal{M}_t} \|S^{t-1}(B^{t-1}(x)) - S(B(x))\|_2$$

4.3 Memory Selection Rule

Since we only store a small number of examples as a way to balance the replay as well as the extra memory cost and training time, we need to carefully select them in order to utilize the memory buffer \mathcal{M} efficiently. Considering that if two stored

examples are very similar, then only storing one of them could possibly achieve similar results in the future. Thus, those stored examples should be as diverse and representative as possible. To this end, after training on t -th task, we employ K-means (MacQueen et al., 1967) to cluster all the examples from current training set S_t : For each $x \in S_t$, we utilize its embedding $B(x)$ as its input feature to conduct K-means. We set the numbers of clusters to $\gamma|S_t|$ and only select the example closest to each cluster’s centroid, following (Wang et al., 2019; Han et al., 2020).

4.4 Overall Objective

We can write the final objective for continual learning on text classification as the following:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{cls} + \mathcal{L}_{nsp} + \mathcal{L}_{task} \\ & + \lambda^g \mathcal{L}_{reg}^g + \lambda^s \mathcal{L}_{reg}^s \end{aligned} \quad (1)$$

We set the coefficient of the first three loss terms to 1 for simplicity and only introduce two coefficients to tune: λ^g and λ^s . In practice, \mathcal{L}_{task} and \mathcal{L}_{cls} are also conducted on each generated *NotNext* example \hat{x} , \mathcal{L}_{reg}^g and \mathcal{L}_{reg}^s are only optimized starting from the second task. The full ID based regularization (IDBR) algorithm is shown in Algorithm 1.

5 Experiment

5.1 Datasets

Following MBPA++ (d’Autume et al., 2019), we use five text classification datasets (Zhang et al., 2015) to evaluate our methods, including AG News (news classification), Yelp (sentiment analysis), DBPedia (Wikipedia article classification), Amazon (sentiment analysis), and Yahoo! Answer (Q&A classification). A summary of the datasets is shown in Table 1. We merge the label space of Amazon and Yelp considering their domain similarity, therefore, there are 33 classes in total.

5.2 Experiment Setup

Due to the limitation of resources, for most of our experiments, we randomly sample 2000 examples per class for every task and create a reduced dataset. See Table 1 for the train/test size of each dataset. We name this setting Setting (Sampled). We tune all the hyperparameters on the basis of Setting (Sampled). Beyond that, to have a comparison with previous State-of-the-art, we also conduct experiments on the same training set and test set as MbPA++ (d’Autume et al., 2019) and LAMOL

Algorithm 1 IDBR

Input Training sets $\{S_1, \dots, S_n\}$, Replay Frequency β , Store ratio γ , Coefficients λ_g, λ_s

Output Optimal models $B, G, S, f_{nsp}, f_{task}, f_{cls}$

```

 $\mathcal{M} = \{\}$   $\triangleright$  Initialize memory buffer
Initialize  $B$  using pretrained BERT
Initialize  $G, S, f_{nsp}, f_{task}, f_{cls}$ 
for  $t = 1, \dots, n$  do
  if  $t \geq 2$  then
    Store  $G(B(x)), S(B(x)), \forall x \in S_t \cup \mathcal{M}$ 
    for batches  $\in S_t$  do
      Optimize  $\mathcal{L}$  in Equation 1
      if step mod  $\beta = 0$  then  $\triangleright$  Replay
        Sample  $t - 1$  batches from  $\mathcal{M}$ 
        Optimize  $\mathcal{L}$  in Equation 1
      end if
    end for
  else  $\triangleright$  No regularization on 1st task
    for batches  $\in S_t$  do
      Optimize  $\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{nsp} + \mathcal{L}_{task}$ 
    end for
  end if
   $\mathcal{C} = \text{K-Means}(S_t, n_{clu}=\gamma|S_t|) \triangleright \mathcal{C} : \text{centroid}$ 
   $\mathcal{C}' = \{ \text{Examples closest to centers} \in \mathcal{C} \}$ 
   $\mathcal{M} \leftarrow \mathcal{M} \cup \mathcal{C}'$   $\triangleright$  Add to memory
end for
return  $B, G, S, f_{nsp}, f_{task}, f_{cls}$ 

```

(Sun et al., 2019), which contains 115,000 training examples and 7,600 test examples for each task. We name the latter Setting (Full).

Our experiments are mainly conducted on the task sequences shown in Table 2. To minimize the effect of task order and task sequence length on the results, we examine both length-3 task sequences and length-5 task sequences in various orders. The first 3 task sequences are a cyclic permutation of $\text{ag} \rightarrow \text{yelp} \rightarrow \text{yahoo}$, which are three classification tasks in different domains (news classification, sentiment analysis, Q&A classification). The last four length-5 task sequences follows d’Autume et al. (2019).

5.3 Baselines

We compare our proposed model with the following baselines in our experiments:

- **Finetune** (Yogatama et al., 2019): fine-tune BERT model sequentially without the episodic memory module and any other loss.
- **Replay** (Wang et al., 2019; d’Autume et al.,

Dataset	Class	Type	Train	Test
AGNews	4	News	8000	7600
Yelp	5	Sentiment	10000	7600
Amazon	5	Sentiment	10000	7600
DBPedia	14	Wikipedia	28000	7600
Yahoo	10	Q&A	20000	7600

Table 1: Dataset statistics we used for Setting (Sampled). Type means the domain of task classification.

Order	Task Sequence
1	ag → yelp → yahoo
2	yelp → yahoo → ag
3	yahoo → ag → yelp
4	ag → yelp → amazon → yahoo → dbpedia
5	yelp → yahoo → amazon → dbpedia → ag
6	dbpedia → yahoo → ag → amazon → yelp
7	yelp → ag → dbpedia → amazon → yahoo

Table 2: Seven random different task sequences used for experiments. The first 6 are used in Setting (Sampled). The last 4 are used in Setting (Full).

2019): **Finetune** model augmented with an episodic memory.

- **Regularization:** On top of **Replay**, with an L2 regularization term.
- **MBPA++** (d’Autume et al., 2019): augment BERT model with an episodic memory module and store **all** seen examples. MBPA++ performs experience replay at training time, and uses KNN to select examples for local adaptation at test time.
- **LAMOL** (Sun et al., 2019): a language model which can generate pseudo training samples used for experience replay. The text classification is performed in Q&A formats.
- **Multi-task learning** (MTL): The model is trained on all tasks simultaneously, which can be considered as an upper-bound for continual learning methods since it has access to data from all tasks at the same time.

5.4 Implementation Details

We use pretrained BERT-based-uncased from HuggingFace Transformers (Wolf et al., 2020) as our base feature extractor. The task generic encoder and task specific encoder are both one linear layer followed by activation function $Tanh$, their output size are both 128 dimensions. The predictors

built on encoders are all one linear layer followed by activation function $softmax$. We use AdamW (Loshchilov and Hutter, 2017) as optimizer. For all modules except the task id predictor, we set the learning rate $lr = 3e^{-5}$; for task id predictor, we set its learning rate $lr_{task} = 5e^{-4}$. The weight decay for all parameters are 0.01.

For experience replay, we set the store ratio $\gamma = 0.01$, i.e. we store 1% of seen examples into the episodic memory module. Besides, we set the replay frequency $\beta = 10$, which means we do experience replay once every ten steps.

For information disentanglement, we mainly tune the coefficients of the regularization loss. For batches from memory buffer \mathcal{M} , we set λ^g to 5.0, select best λ^s from $\{1.0, 2.0, 3.0, 4.0, 5.0\}$. For batches from current training set S , we set λ^g to 0.5, select best λ^s from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$.

6 Results and Discussion

We evaluated models after training on all tasks and reported their average accuracies on all test sets as our metric. Table 3 summarizes our results in Setting (Sampled). While continual finetuning suffered from severe forgetting, experience replay with 1% stored examples achieved promising results, which demonstrated the importance of experience replay for continual learning in NLP. Beyond that, simple regularization turned out to be a robust method on the basis of experience replay, which showed consistent improvements on all 6 orders. Our proposed Information Disentanglement Based Regularization (IDBR) further improves regularization consistently under all circumstances.

Table 4 compares IDBR with previous SOTA: MBPA++ and LAMOL in Setting (Full). We list some differences among our settings in the table. Although MBPA++ applies local adaptation when testing, IDBR still outperformed it by an obvious margin. We achieved comparative results with LAMOL even if they need task identifiers during inference, which makes their prediction easier.

6.1 Impact of the Lengths of Task Sequences

Comparing results of length-3 sequences and length-5 sequences in Table 3, we found that the gap between IDBR and multi-task learning became bigger when the length of task sequence changed from 3 to 5. To better understand how IDBR gradually forgot, we followed Chaudhry et al. (2018a)

Model	Length-3 Task Sequences				Length-5 Task Sequences			
Order	1	2	3	Average	4	5	6	Average
Finetune	25.79	36.56	41.01	34.45	32.37	32.22	26.44	30.34
Replay	69.32	70.25	71.31	70.29	68.25	70.52	70.24	69.67
Regularization	71.50	70.88	72.93	71.77	72.28	73.03	72.92	72.74
IDBR	71.80	72.72	73.08	72.53	72.63	73.72	73.23	73.19
MTL	74.16	74.16	74.16	74.16	75.09	75.09	75.09	75.09

Table 3: Summary of results on Setting (Sampled) using averaged accuracy after training on the last task. All results are averaged over 3 runs. The p -values of paired t -test between nine numbers of Regularization and IDBR are 0.018 on Length-3 and 0.009 on Length-5, demonstrating the significant differences.

Model	TT	TI	LA	LM	Length-5 Task Sequences				
Order					4	5	6	7	Average
MBPA++ †			✓		70.7	70.2	70.9	70.8	70.7
MBPA++ ††			✓		74.9	73.1	74.9	74.1	74.3
LAMOL ††	✓	✓		✓	76.1	76.1	77.2	76.7	76.5
IDBR	✓				75.9	76.2	76.4	76.7	76.3

Table 4: Summary of results on Setting (Full) using averaged accuracy after training on the last task. Our results are averaged over 3 runs. † means we fetch numbers from d’Autume et al. (2019). †† means we fetch numbers from Sun et al. (2019). **TT**: whether task-id is available during training. **TI**: whether task-id is available during inference. **LA**: whether need local adaptation during inference. **LM**: whether train a language model.

Order	4	5	6	Average
After 2 tasks	0.64	1.63	0.07	0.78
After 3 tasks	3.18	2.56	1.56	2.43
After 4 tasks	3.60	2.17	2.20	2.66
After 5 tasks	3.46	2.33	2.88	2.89

Table 5: Forgetting measure Chaudhry et al. (2018a) calculated every time finish training on a new task. All results are averaged over 3 runs.

to measure *forgetting* as follows:

$$\mathcal{F}_k = \mathbb{E}_{j=1\dots t-1} f_j^k,$$

$$f_j^k = \max_{l \in \{1\dots k-1\}} a_{l,j} - a_{k,j}$$

where $a_{l,j}$ is the model’s accuracy on task j after trained on task l . On order 4, 5 and 6, we calculated the forgetting every time after IDBR was trained on a new task and summarized them in Table 5. For continual learning, we hypothesize that the model is prone to suffer from more severe forgetting as the task sequence becomes longer. We found that although there was some big drop after training on the 3rd task, IDBR maintained stable performance as the length of task sequence increased, especially after training on 4-th and 5-th task, the forgetting increment was relatively small, which demonstrated the robustness of IDBR.

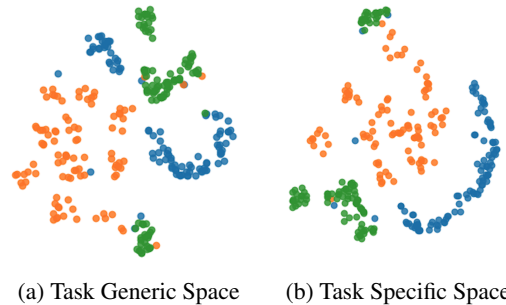


Figure 2: t-SNE visualization of task generic hidden space and task specific hidden space of IDBR

6.2 Visualizing Disentangled Spaces

To study whether our task generic encoder G tends to learn more generic information and task specific encoder S captures more task specific information, we used t-SNE (Maaten and Hinton, 2008) to visualize the two hidden spaces of IDBR, using the final model trained on order 2, and the results are shown in Figure 2, where Figure 2a visualizes task generic space and Figure 2b visualizes task specific space. We observe that compared with task specific space, generic features from different tasks were more mixed, which demonstrates that the next sentence prediction helped task generic space to be more task-agnostic than task specific space, which was induced to learn separated representations for

Model	Accuracy
Regularization	73.03
IDBR w/o \mathcal{L}_{nsp}	73.17
IDBR w/o \mathcal{L}_{tak}	73.29
IDBR	73.72

Table 6: Comparison among using task-id prediction only, next sentence prediction only and both of them. All results are averaged over 3 runs.

Model	4	5	6	Avg
Reg only on s	72.05	72.54	72.61	72.40
Reg only on g	72.01	72.98	72.73	72.57
Reg on both	72.63	73.72	73.23	73.19

Table 7: Comparison among using regularization on task specific space only, task generic space only and both of them. All results are averaged over 3 runs.

different tasks. Considering we only employed two simple auxiliary tasks, the effect of information disentanglement was noticeable.

6.3 Ablation Studies

Effect of Disentanglement In order to demonstrate that each module of our information disentanglement helps the learning process, we performed ablation study on the two auxiliary tasks using order 5 as a case study. The results are summarized in Table 6. We found that both task-id prediction and next sentence prediction contribute to the final performance. Furthermore, the performance gain was much larger by combing these two auxiliary tasks together. Intuitively, the model needs both tasks to disentangle the representation well, since it is easy for the model to ignore one of the spaces if the constraint is not imposed appropriately. The results show that the two tasks are likely complementary to each other in helping the model learn better disentangled representations.

Impact of Regularization To study the effect of regularization on task generic hidden space g and task specific hidden space s , we performed an ablation study which only applied regularization on g or s , and compared the results with regularization on both in Table 7. We found that regularization on both spaces results in a much better performance than regularization only on one of them only, which demonstrates the necessity of both regularizers. While we may expect to give more tolerance to specific space for changing, we found that no regularization on it would lead to severe forgetting

Rules	1	2	3	Average
Random	71.52	72.60	73.03	72.38
K-Means	71.80	72.72	73.05	72.52

Table 8: Comparison between different selection rules: select stored examples randomly or select by K-Means. All results are averaged over 3 runs.

of previously learnt good task specific embeddings, hence it is necessary to add a regularizer over this space as well. Beyond that, we also observed that under most circumstances, adding regularization on the task generic space g results in a more significant gain than adding regularization on the task specific space s , consistent with our intuition that task generic space changes less across tasks and thus preserving it better helps more in alleviating catastrophic forgetting.

Impact of K-Means To demonstrate our hypothesis that when the memory budget is limited, selecting the most representative subset of examples is vital to the success of continual learning, we performed an ablation study on order 1,2,3 using IDBR with and without K-Means. The result is shown in Table 8. From the table, we found that using K-Means helps boost the overall performance to some extent. Specifically, the improvement brought by K-Means was larger on those challenging orders, i.e. orders on which IDBR had worse performance. This is because for these challenging orders, the forgetting is more severe and the model needs more examples from previous tasks to help it retain previous knowledge. Thus with the same memory budget constraint, diversity across saved examples will help the model better recover knowledge learned from previous tasks.

7 Conclusion

In this work, we introduce an information disentanglement based regularization (IDBR) method for continual text classification, where we disentangle the hidden space into task generic space and task specific space and further regularize them differently. We also leverage K-Means as the memory selection rule to help the model benefit from the augmented episodic memory module. Experiments conducted on five benchmark datasets demonstrate that IDBR achieves better performances compared to previous state-of-the-art baselines on sequences of text classification tasks with various orders and lengths. We believe the proposed approach can be

extended to continual learning for other NLP tasks such as sequence generation and sequence labeling as well, and plan to explore them in the future.

Acknowledgment

We would like to thank the anonymous reviewers for their helpful comments, and the members of Georgia Tech SALT group for their 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.
- Yu Bao, Hao Zhou, Shujian Huang, Lei Li, Lili Mou, Olga Vechtomova, Xinyu Dai, and Jiajun Chen. 2019. Generating sentences from disentangled syntactic and semantic spaces. *arXiv preprint arXiv:1907.05789*.
- Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. 2018a. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–547.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2018b. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. 2019. Episodic memory in lifelong language learning. *arXiv preprint arXiv:1906.01076*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Zhenxin Fu, Xiaoye Tan, Nanyun Peng, Dongyan Zhao, and Rui Yan. 2017. Style transfer in text: Exploration and evaluation. *arXiv preprint arXiv:1711.06861*.
- 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.
- Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. Meta-learning with sparse experience replay for lifelong language learning. *arXiv preprint arXiv:2009.04891*.
- Khurram Javed and Martha White. 2019. Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, pages 1820–1830.
- Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2018. Disentangled representation learning for non-parallel text style transfer. *arXiv preprint arXiv:1808.04339*.
- 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.
- Yuan Li, Chunyuan Li, Yizhe Zhang, Xiujun Li, Guoqing Zheng, Lawrence Carin, and Jianfeng Gao. 2020. Complementary auxiliary classifiers for label-conditional text generation. In *AAAI*, pages 8303–8310.
- Zhizhong Li and Derek Hoiem. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, pages 6467–6476.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- James MacQueen et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA.
- Arun Mallya and Svetlana Lazebnik. 2018. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Abiola Obamuyide and Andreas Vlachos. 2019. Meta-learning improves lifelong relation extraction. In *Proceedings of the 4th Workshop on Representation Learning for NLP (Repl4NLP-2019)*, pages 224–229.

- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71.
- David N Perkins, Gavriel Salomon, et al. 1992. Transfer of learning. *International encyclopedia of education*, 2:6452–6457.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Amal Rannen, Rahaf Aljundi, Matthew B Blaschko, and Tinne Tuytelaars. 2017. Encoder based lifelong learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1320–1328.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010.
- Mark B Ring. 1998. Child: A first step towards continual learning. In *Learning to learn*, pages 261–292. Springer.
- Alexey Romanov, Anna Rumshisky, Anna Rogers, and David Donahue. 2018. Adversarial decomposition of text representation. *arXiv preprint arXiv:1808.09042*.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2019. Lamol: Language modeling for lifelong language learning. In *International Conference on Learning Representations*.
- Sebastian Thrun. 1998. Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer.
- Hong Wang, Wenhan Xiong, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019. Sentence embedding alignment for lifelong relation extraction. *arXiv preprint arXiv:1903.02588*.
- Zirui Wang, Sanket Vaibhav Mehta, Barnabás Póczos, and Jaime Carbonell. 2020. Efficient meta lifelong-learning with limited memory. *arXiv preprint arXiv:2010.02500*.
- Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45.
- Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, et al. 2019. Learning and evaluating general linguistic intelligence. *arXiv preprint arXiv:1901.11373*.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. *Proceedings of machine learning research*, 70:3987.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657.
- Junbo Zhao, Yoon Kim, Kelly Zhang, Alexander Rush, and Yann LeCun. 2018. Adversarially regularized autoencoders. In *International Conference on Machine Learning*, pages 5902–5911. PMLR.