

Continual Training of Language Models for Few-Shot Learning

Zixuan Ke¹, Haowei Lin², Yijia Shao², Hu Xu¹, Lei Shu^{1*} and Bing Liu¹

¹Department of Computer Science, University of Illinois at Chicago

²Wangxuan Institute of Computer Technology, Peking University

¹{zke4, hxu48, liub}@uic.edu

²{linhaowei, shaoyj}@pku.edu.cn

Abstract

Recent work on applying large language models (LMs) achieves impressive performance in many NLP applications. Adapting or post-training an LM using an unlabeled domain corpus can produce even better performance for end-tasks in the domain. This paper proposes the problem of continually extending an LM by incrementally post-train the LM with a sequence of unlabeled domain corpora to expand its knowledge without forgetting its previous skills. The goal is to improve the few-shot end-task learning in these domains. The resulting system is called CPT (Continual Post-Training), which to our knowledge, is the first continual post-training system. Experimental results verify its effectiveness.¹

1 Introduction

Recent work has shown that large LMs have the ability to perform few-shot (or even zero-shot) learning well (Brown et al., 2020b; Rae et al., 2021; Smith et al., 2022). *Post-training* (a.k.a., *domain-adaptive pre-training* or *pre-finetuning*) an LM with a large unlabeled domain corpus before end-task fine-tuning in the domain achieves better results (Xu et al., 2019; Gururangan et al., 2020a) than directly fine-tuning the LM. This paper goes a step further to study the problem of improving an LM’s ability to handle new and ever emerging domains. For this, one needs to *continually post-train* the LM with a sequence of domains. A key issue associated with this problem is *catastrophic forgetting* (CF).² This paper thus investigates how to continually extend the LM’s knowledge without suffering from CF. From a broader perspective, since training a large LM from scratch is extremely

expensive and computation intensive, incrementally updating the LM with the latest language data reflecting the ever changing development of the language itself, social events and the knowledge from different fields is becoming more and more critical. As humans are very effective at incremental learning, if we can imitate this human capability with little or no forgetting, we will be pushing the AI research forward significantly.

The proposed system, called CPT, is a continual learning (CL) system for post-training. Starting from a pre-trained LM (e.g., RoBERTa (Liu et al., 2019b)), it incrementally post-trains the LM with a sequence of domains using their unlabeled corpora. Once a task (a domain in our case)³ is trained, its data is no longer accessible. At any time, the resulting continually post-trained LM can be used by end-tasks in the trained domains. This is in the *task-incremental learning* (TIL) setting of CL, where the task id (domain id in our case) is provided when the learned model of a task needs to be used later (the use of domain id is discussed in Sec. 2.1).⁴ This paper proposes an effective approach called CPT and focuses on the challenging and practical scenario of *few-shot* end-task learning after post-training a sequence of domains.

Continual post-training is different from conventional CL (Chen and Liu, 2018). The key difference is that in conventional CL, each task is an end-task, but in our case the end-task involves fine-tuning the continual post-trained LM (called p-LM). This causes major forgetting, which we call the *catastrophic butterfly effect* (CBE) and does not happen in conventional CL. Our proposed system, CPT, can solve both CF and CBE, based on a novel hard masking mechanism (Sec. 2.2) and can achieve *no* forgetting. As shown in Sec. 3.3, naively ap-

*Now at Google Research leishu@google.com

¹<https://github.com/UIC-Liu-Lab/CPT>

²CF means that learning a new task/domain may need to modify the existing network, which degrades the performance of previous tasks/domains (McCloskey and Cohen, 1989).

³We will use the term *domain* in this paper to be consistent with the post-training literature

⁴CL has two other settings: *class-incremental learning* and *domain-incremental learning* (van de Ven and Tolias, 2019).

plied existing CL systems cannot effectively prevent CF (even though some existing techniques have shown almost perfect CF prevention ability in conventional CL).

Experiments in 4 domains and their corresponding end-tasks demonstrate the effectiveness of the proposed CPT system.

Related Work. Overcoming CF is a major goal of CL (Chen and Liu, 2018). There are many existing approaches, e.g., *regularization-based approaches* (Kirkpatrick et al., 2016; Seff et al., 2017), *replay-based approaches* (Rebuffi et al., 2017; Lopez-Paz and Ranzato, 2017) and *parameter isolation based approaches* (Serrà et al., 2018; Fernando et al., 2017). Our CPT is based on parameter isolation and uses masks in continual post-training. Recently, CL has drawn attention in NLP. It has been used for slot filling (Shen et al., 2019), language learning (Li et al., 2019), sentence embedding (Liu et al., 2019a), translation (Khayrallah et al., 2018), cross-lingual modeling (Liu et al., 2020b), question answering (Greco et al., 2019) and text classification (Ke et al., 2021a,b; Sun et al., 2020; Huang et al., 2021; Chuang et al., 2020; Mehta et al., 2021; Madotto et al., 2020). However, none of them tries to improve an LM.

CPT is closely related to ELLE (Qin et al., 2022), which does *continual pre-training*. The key difference is that ELLE starts from random initialization, while our CPT starts from a pre-trained LM. We tried to adapt ELLE for continual post-training by learning from a pre-trained RoBERTa but it fails to converge. This also indicates it is non-trivial to do well in our setting. Readers can refer to Appendix A for a full coverage of the related work.

2 Proposed CPT System

CPT continually post-trains RoBERTa (Liu et al., 2019b). This is achieved by two *continual learning plug-in* (called CL-plugin) modules inserted into each transformer layer of RoBERTa. CL-plugin is inspired by adapters in (Houlsby et al., 2019). While adapters can isolate different tasks, one needs to allocate a new adapter for each task and no knowledge can be shared among different tasks’ adapters. The CL-plugin, however, is a CL system that learns a sequence of tasks with adapters shared by all domains. Figure 1 gives the CPT architecture with two CL-plugins added to RoBERTa.

Sequential vs. Parallel CL-plugin. Instead of following the original sequential adapter (Houlsby

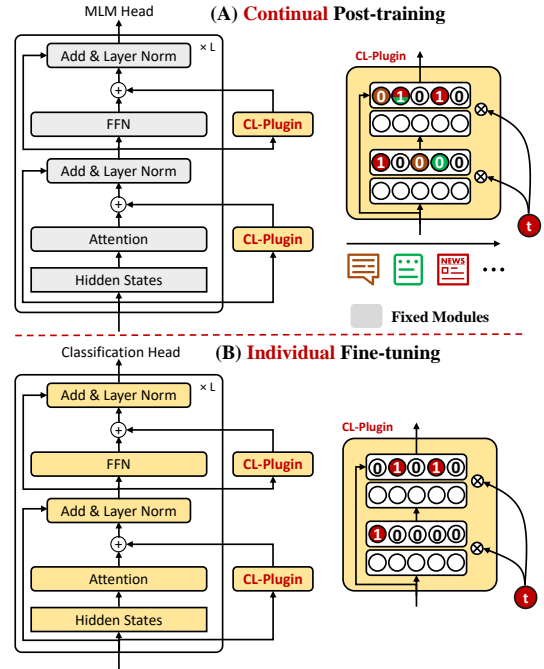


Figure 1: Architecture of CPT, which has two CL-plugins inserted in the transformer layers of RoBERTa in a parallel manner (FFN: feed-forward network). (A) CPT for *continual post-training*. It uses a masked language model (MLM) head for unsupervised post-training of the CL-plugins only. (B) CPT for *individual fine-tuning*. CPT is evaluated by the corresponding individual end-task performance of all post-trained tasks. Each CL-plugin has numbers and colors indicating its masks and is illustrated in Appendix B.

et al., 2019), CL-plugin adopts the parallel adapter idea in (He et al., 2021). The difference is that the former inserts an adapter *after* the FFN/attention layer while the latter inserts it *before* FFN/attention layer (see Fig. 1). We choose the parallel version as it performs better (see Sec. 3.3).

In post-training, only the two CL-plugins are trained. The components of the original pre-trained RoBERTa are fixed. In end-task fine-tuning, all components are trainable. A CL-plugin is a two-layer fully connected network with a task mask mechanism. It takes two inputs: (1) hidden states $h^{(t)}$ from the feed-forward layer in a transformer layer and (2) task ID t needed by task incremental learning (TIL). Inside a CL-plugin, task masks (TMs), which indicate task-specific neurons, are used to deal with CF. Since TMs is differentiable, the whole CPT can be trained end-to-end.

2.1 Task Masks (TMs)

In each layer of a CL-plugin, task masks are used to protect those neurons that are important for pre-

vious tasks to overcome CF. The masks basically forbid gradient updates to those neurons during backpropagation in learning a new task. Note that a task is also a domain in our case.

Learning a new task/domain consists of two main steps: (1) apply the mask in each layer for each old task to block off the gradient flow to protect the model for the old task, and (2) learn domain t and its masks for future use. We present (2) first.

Learning Task Masks for Overcoming CF. In learning each task t , a mask (a “soft” binary mask) $m_l^{(t)}$ is trained for the task at each layer l in CL-plugin, indicating the neurons that are important for the task. We borrow the hard attention idea in (Serrà et al., 2018) and leverage the task ID embedding to train the mask. For a task ID t , its embedding $e_l^{(t)}$ consists of differentiable deterministic parameters that can be learned together with other parts of the network. To generate the task mask $m_l^{(t)}$ from $e_l^{(t)}$, *Sigmoid* is used as a pseudo-gate (mask) function. $m_l^{(t)}$ is computed with

$$m_l^{(t)} = \sigma(e_l^{(t)} / \tau), \quad (1)$$

where τ is a temperature variable, linearly annealed from 1 to τ_{\min} (a small positive value).

In the forward pass, given the output of each layer l , $k_l^{(t)}$, we element-wise multiply mask $m_l^{(t)}$,

$$o_l^{(t)} = k_l^{(t)} \otimes m_l^{(t)}. \quad (2)$$

The masked output $o_l^{(t)}$ of the last layer in CL-plugin is fed to the next layer of the RoBERTa with a skip-connection. After learning task t , the final $m_l^{(t)}$ is saved and added to the set $\{m_l^{(t)}\}$.

Applying Task Masks. Before learning a new task t , we first accumulate and set the masks $m_l^{(i_{\text{prev}})}$ on the neurons in each layer l for all old tasks i_{prev} so that in backpropagation, the gradient $g_l^{(t)}$ for task t will not flow to these neurons. Since $m_l^{(i_{\text{prev}})}$ is pseudo binary, we use max-pooling to achieve the accumulation and condition the gradient:

$$g_l^{(t)} = g_l^{(t)} \otimes (1 - (\text{MaxPool}(\{m_l^{(i_{\text{prev}})}\}))). \quad (3)$$

Those gradients corresponding to the 1 entries in $\text{MaxPool}(\{m_l^{(i_{\text{prev}})}\})$ are set to 0 (to block off gradient flow) while the others remain unchanged. In this way, neurons in old tasks are protected. Note that we expand (copy) the vector $m_l^{(t_a)}$ to match the dimensions of $g_l^{(t)}$.

2.2 Catastrophic Butterfly Effect in Fine-tuning

To perform an end-task in a post-trained domain, we fine-tune the mask-protected model of the domain, which is indicated by the task/domain id. The fine-tuning uses the corresponding domain neurons for the specific end-task by setting $\tau = \tau_{\min}$ and condition the output via Eq. 2. With the masks, there should be no forgetting for continual post-training and the end-task fine-tuning performance should be similar to post-train each domain separately. However, we found that this is not the case.⁵ Our investigation found that the problem is due to the *pseudo-gate* function in Eq. 1. No matter how small τ is, Eq. 1 can only give us a mask almost 0 (or 1). This causes the following: (1) During post-training, the gradients for used neurons in Eq. 3 are not exactly 0 but a very small number. (2) During fine-tuning, we cannot make use of the corresponding neurons for the specific end-task by simply setting $\tau = \tau_{\min}$. The small change in the neurons for old domains during post-training caused by (1) is neglect-able in conventional CL because in conventional CL we evaluate the model using test sets and no weights update involved. However, in CPT, the end-task needs to fine-tune the continually post-trained LM model (p-LM), which involves weight updating. A small change to the p-LM during continual post-training can result in a different initialization for the end-task fine-tuning and give totally different fine-tuning results. We call this *butterfly effect* inspired by the term indicating a small state change in nature (e.g., the flap of a butterfly’s wings in Brazil) can result in large differences in a later state (e.g., a tornado in Texas).

We propose a simple method to solve it, i.e., adding a threshold θ to the $m_l^{(t)}$ to make it a *hard binary mask*,

$$m_l^{(t)} = \begin{cases} 1, & m_l^{(t)} > \theta, \\ 0, & m_l^{(t)} < \theta. \end{cases} \quad (4)$$

We then apply it to Eq. 3 in gradient manipulation and Eq. 2 in end-task fine-tuning. θ can be easily set (we use 0.5) since Eq. 1 already gives a pseudo-binary mask. Note that this has almost no effect on post-training as it is used to block the backward

⁵For example, fine-tuning an end restaurant sentiment classification task achieves macro-F1 (MF1) of 0.64 right after post-training the restaurant domain but its fine-tuning MF1 drops to 0.44 after post-training three more domains.

pass gradient flow during post-training and select the corresponding neurons during fine-tuning.

3 Experiments

The proposed paradigm uses a different evaluation from that of conventional continual learning (CL). After unsupervised continual post-training of an LM (RoBERTa in our case) with a sequence of domains, the resulting p-LM is used to fine-tune an end few-shot classification task from any post-trained domain. There is no CL during end-task fine-tuning. Each fine-tuning task is done separately.

3.1 Datasets and Baselines

Datasets: We use 4 unlabeled domain datasets: *Yelp Restaurant* (Xu et al., 2019), *AI Papers* (Lo et al., 2020), *ACL Papers* (Lo et al., 2020) and *AG-News* (Zhang et al., 2015) and their 4 corresponding end-task classification datasets.⁶

Baselines. Since no existing method can perform our task, we use 6 non-CL and 7 adapted CL methods as our baselines. The non-CL baselines include (1) **RoBERTa** and (2) **Adapter** where we directly fine-tune the pre-trained model or adapter (without any post-training); (3) **RoBERTa-ONE**, (4) **Adapter-ONE** and (5) **Prompt-ONE**, where we build a model for each task using a separate network. It has no knowledge transfer or CF. (6) **DEMIX** (Gururangan et al., 2021) trains a separate adapter for each task and initializes the adapter from its most similar previous task adapter. The 7 adapted CL baselines include (7) **RoBERTa-NCL** and (8) **Adapter-NCL**, where we post-train the domains one by one with no mechanism to deal with CF/transfer. Other are state-of-the-art CL baselines and we adapt them for continual post-training.⁷

3.2 Implementation Details

Architecture. We adopt RoBERTa_{BASE} as our backbone LM. A masked language model head is applied for post-training. The fine-tuning follows the standard practice (Devlin et al., 2019), where we pass the final layer $\langle /s \rangle$ token representation

⁶These are popularly used in related works. Details of the datasets are given in Appendix C. We conduct experiments using *few-shot* learning end-tasks. Following (Gu et al., 2021), we use 32 training samples for *Restaurant* and *AGNews*, 48 training samples for *ACL* and 56 training samples for *AI* due to different numbers of classes in each dataset.

⁷Readers can refer to Appendix D for the detailed of each of these baselines.

to a task-specific feed-forward layer for prediction. The feed-forward layer with softmax output is used as the classification heads, together with the categorical cross-entropy loss. Note that for the aspect sentiment classification task (see Table 3), we adopt the ASC formulation in (Xu et al., 2019), where the aspect (e.g., “*sound*”) and review sentence (e.g., “*The sound is great*”) are concatenated via $\langle /s \rangle$.

Hyperparameters. Unless otherwise stated, the same hyper-parameters are used in all experiments. We use 0.0025 for τ_{\min} in Eq. 1 and θ is set to 0.5 in Eq. 4 in the main paper. As shown in Figure 1, there are two CL-plugins for each Transformer layer (one at the bottom in parallel with attention and one at the top in parallel with FFN). We search the CL-plugin size within {128, 256, 512, 768, 1024} and adopt 512 for the bottom one and 768 for the top one based on validation experiments. The task id embeddings have the same size as the hidden layer dimension of the CL-plugin. The maximum input length is set to 164 which is long enough for all datasets. We use Adam optimizer and set the learning rate to $1e-4$ for post-training and $5e-5$ for fine-tuning. The batch size is set to 48 for post-training and 20 for fine-tuning. Since each of our domain-specific dataset has a different size, we train CPT on each task/domain for 1 epoch for post-training, which is approximately 13K steps, following (Gururangan et al., 2020b; Xu et al., 2019). We train on end-task fine-tuning datasets for 20 epochs and take the results for the last epoch, assuming no validation sets. We empirically found 20 epochs can give us a relatively stable results.

3.3 Evaluation Results and Analysis

We report the average results of the 4 different fine-tuning tasks (or datasets) in accuracy and Macro-F1 after post-training on all unlabeled domain datasets in Table 1. The forgetting rate (forget R.) (Liu et al., 2020a) is also reported. The higher the forgetting rate is, the more forgetting it has. Negative rates indicate positive knowledge transfer.⁸

Superiority of CPT. Clearly, CPT outperforms all baselines and achieves no forgetting. More specifically, CPT markedly outperforms the two

⁸Forgetting rate is computed as follows (Liu et al., 2020a), $\frac{1}{t-1} \sum_{i=1}^{t-1} A_{i,i} - A_{t,i}$, where $A_{i,i}$ is the end-task performance right after its domain i is post-trained, and $A_{t,i}$ is the performance of the end-task of domain i after post-training the last domain. We average over all end-tasks except the last one as the last domain has no forgetting.

Category	Domain Model	Restaurant		AI		ACL		AGNews		Average		Forget R.	
		MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc
Non-CL	RoBERTa	50.61	74.77	27.88	28.44	32.19	34.59	64.19	65.95	43.72	50.94	—	—
	Adapter	45.40	67.28	23.69	24.56	24.99	27.55	64.53	66.50	39.65	46.48	—	—
	RoBERTa-ONE	53.63	76.73	29.86	30.11	33.05	35.72	62.57	65.13	44.78	51.92	—	—
	Adapter-ONE	52.19	74.20	30.80	31.59	36.59	36.99	61.66	63.94	45.31	51.68	—	—
	Prompt-ONE	28.93	59.79	21.06	22.10	28.02	29.22	60.70	62.58	34.68	43.42	—	—
	DEMIX	53.14	75.28	27.68	27.29	37.63	38.57	63.18	65.13	45.41	51.57	—	—
CL	RoBERTa-NCL	42.59	67.56	31.57	31.62	33.07	34.54	60.18	63.50	41.85	49.30	3.27	2.82
	Adapter-NCL	47.42	70.23	29.56	29.90	35.92	37.58	61.73	64.45	43.65	50.54	2.21	1.69
	HAT	50.45	71.78	28.33	29.41	34.93	37.15	62.97	65.05	44.17	50.85	2.43	2.04
	BCL	51.70	74.34	29.66	30.96	32.85	34.82	63.60	65.47	44.45	51.40	1.47	0.82
	KD	39.75	67.11	29.63	29.33	38.30	42.09	62.85	65.39	42.63	50.98	4.92	3.07
	EWC	48.32	71.59	30.96	31.01	35.96	38.05	62.29	64.95	44.38	51.40	1.40	0.80
	DER++	48.09	71.79	30.71	30.54	34.25	35.77	64.24	66.11	44.32	51.05	1.79	1.62
	CPT	53.90	75.13	30.42	30.89	37.56	38.53	63.77	65.79	46.41	52.59	0.00	0.00

Table 1: End-task macro-F1 (MF1), accuracy and forgetting rate results for all domains *after continual post-training of all domains*. The results are averages of 5 random seeds (the domain training order is as they appear in the first row). Due to space limits, the results for *different domain orders* and the *standard deviations* are reported in Appendix E and Appendix F, respectively). Non-CL baselines has no forgetting.

baselines without post-training (RoBERTa and Adapter), indicating CPT can learn new domains well. These two baselines are also significantly worse than other baselines, indicating that fine-tuning the pre-trained RoBERTa alone is weak. Comparing with CL baselines, **CPT achieves no forgetting** (we can see the forgetting rate is 0), indicating the high effectiveness of the proposed approach. We also note that CPT is even slightly better than those ONE baselines, indicating *some positive knowledge transfer in CPT*.

3.4 Ablations

In Table 2, we give the ablation results. We are interested in the following:

(1) **Catastrophic butterfly effect (CBE)**. The third row with “w/o butterfly” shows results without the hard binary mask mechanism in Eq. 4. Clearly, the results are worse and the model suffers from forgetting. This indicates CBE and our approach is effective.

(2) **Different Architecture**. CPT is based on CL-plugin, which is inspired by adapters. Another popular way to use adapters is to make it sequential (Houlsby et al., 2019). Sequential adapter (first row) is clearly poorer than our parallel one. This conforms to the observation in (He et al., 2021).

(3) **Different Orders**. Table 1 only reports the results of one fixed domain order (Restaurant→AI→ACL→AGNews). We are interested in how the order impacts CPT results. We give the detailed results for all the other baselines and detailed domain orders in Appendix E. We can see the results of CPT does not change much and it still outperforms other baselines. This indicates the

Model	Final Performance	
	MF1	Acc
CPT (Sequential Adapter)	43.00	50.25
CPT (w/o butterfly)	44.17	50.85
CPT (w/o masking)	43.65	50.54
CPT	46.41	52.58

Table 2: Ablation experiment results.

CPT’s robustness to domain orders in post-training.

4 Conclusion

This paper proposed to continually post-train an LM with a sequence of domains using their unlabeled domain corpora. An effective method (CPT) is also proposed. An end-task from any post-trained domain can fine-tune the resulting LM. Experimental results demonstrate the effectiveness of CPT.

5 Limitations

We list two limitations of CPT. First, CPT adds CL-plugins for continual post-training with no change to the underlying LM in training. Although a CL-plugin is small compared to an LM, it is still interesting and may be more effective to explore the idea of updating the LM directly without any additional modules. Second, domain ids are needed in both training and testing for CPT. In some applications, it may be hard to provide a domain id for each fine-tuning end-task. We will explore these in our future work as specializing and/or incrementally improving an LM is an important problem.

Acknowledgments

The work of Zixuan Ke and Bing Liu was supported in part by three National Science Foundation grants (IIS-1910424, IIS-1838770, and CNS-2225427).

References

- Antreas Antoniou, Massimiliano Patacchiola, Mateusz Ochal, and Amos J. Storkey. 2020. [Defining benchmarks for continual few-shot learning](#). *CoRR*, abs/2004.11967.
- Sagie Benaim and Lior Wolf. 2018. [One-shot unsupervised cross domain translation](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2108–2118.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020a. Language models are few-shot learners. *Advances in neural information processing systems*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020b. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. 2020. Dark experience for general continual learning: a strong, simple baseline. *arXiv preprint arXiv:2004.07211*.
- Zhiyuan Chen and Bing Liu. 2018. Lifelong machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–207.
- Yung-Sung Chuang, Shang-Yu Su, and Yun-Nung Chen. 2020. Lifelong language knowledge distillation. *arXiv preprint arXiv:2010.02123*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. 2017. Pathnet: Evolution channels gradient descent in super neural networks. *CoRR*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 3816–3830. Association for Computational Linguistics.
- Tianyu Gao, Xu Han, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2020. [Neural snowball for few-shot relation learning](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7772–7779. AAAI Press.
- Claudio Greco, Barbara Plank, Raquel Fernández, and Raffaella Bernardi. 2019. Psycholinguistics meets continual learning: Measuring catastrophic forgetting in visual question answering. *arXiv preprint arXiv:1906.04229*.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Yiduo Guo, Bing Liu, and Dongyan Zhao. 2022. Online continual learning through mutual information maximization. In *International Conference on Machine Learning*, pages 8109–8126. PMLR.
- Suchin Gururangan, Mike Lewis, Ari Holtzman, Noah A. Smith, and Luke Zettlemoyer. 2021. [Demix layers: Disentangling domains for modular language modeling](#).
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. 2020a. Don’t stop pretraining: adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*.
- Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020b. Don’t stop pretraining: Adapt language models to domains and tasks. In *ACL*.
- Junxian He, Chunting Zhou, Xueze Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2021. [Towards a unified view of parameter-efficient transfer learning](#). *CoRR*, abs/2110.04366.
- Xu He and Herbert Jaeger. 2018. Overcoming catastrophic interference using conceptor-aided backpropagation. In *ICLR*.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7).
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *ICML*.

- Zikun Hu, Xiang Li, Cunchao Tu, Zhiyuan Liu, and Maosong Sun. 2018. [Few-shot charge prediction with discriminative legal attributes](#). In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018, Santa Fe, New Mexico, USA, August 20-26, 2018*, pages 487–498. Association for Computational Linguistics.
- Yufan Huang, Yanzhe Zhang, Jiaao Chen, Xuezhi Wang, and Diyi Yang. 2021. Continual learning for text classification with information disentanglement based regularization. *arXiv preprint arXiv:2104.05489*.
- Xisen Jin, Bill Yuchen Lin, Mohammad Rostami, and Xiang Ren. 2021. [Learn continually, generalize rapidly: Lifelong knowledge accumulation for few-shot learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 714–729. Association for Computational Linguistics.
- David Jurgens, Srijan Kumar, Raine Hoover, Daniel A. McFarland, and Dan Jurafsky. 2018. Measuring the evolution of a scientific field through citation frames. *TACL*.
- Zixuan Ke, Bing Liu, Nianzu Ma, Hu Xu, and Lei Shu. 2021a. Achieving forgetting prevention and knowledge transfer in continual learning. *NeurIPS*.
- Zixuan Ke, Bing Liu, Hu Xu, and Lei Shu. 2021b. Classic: Continual and contrastive learning of aspect sentiment classification tasks. In *EMNLP*.
- Zixuan Ke, Hu Xu, and Bing Liu. 2021c. Adapting bert for continual learning of a sequence of aspect sentiment classification tasks. In *NAACL*, pages 4746–4755.
- Huda Khayrallah, Brian Thompson, Kevin Duh, and Philipp Koehn. 2018. Regularized training objective for continued training for domain adaptation in neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 36–44.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2016. Overcoming catastrophic forgetting in neural networks. *CoRR*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *EMNLP*.
- Yuanpeng Li, Liang Zhao, Kenneth Church, and Mohamed Elhoseiny. 2019. Compositional language continual learning. In *International Conference on Learning Representations*.
- Tianlin Liu, Lyle Ungar, and Joao Sedoc. 2019a. Continual learning for sentence representations using conceptors. *arXiv preprint arXiv:1904.09187*.
- Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. 2020a. [Mnemonics training: Multi-class incremental learning without forgetting](#). In *CVPR*, pages 12242–12251. IEEE.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized BERT pretraining approach. *CoRR*.
- Zihan Liu, Genta Indra Winata, Andrea Madotto, and Pascale Fung. 2020b. Exploring fine-tuning techniques for pre-trained cross-lingual models via continual learning. *arXiv preprint arXiv:2004.14218*.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel S. Weld. 2020. S2ORC: the semantic scholar open research corpus. In *ACL*.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *NIPS*.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In *ACL*.
- Andrea Madotto, Zhaojiang Lin, Zhenpeng Zhou, Seungwhan Moon, Paul Crook, Bing Liu, Zhou Yu, Eunjoon Cho, and Zhiguang Wang. 2020. Continual learning in task-oriented dialogue systems. *arXiv preprint arXiv:2012.15504*.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*.
- Sanket Vaibhav Mehta, Darshan Patil, Sarath Chandar, and Emma Strubell. 2021. An empirical investigation of the role of pre-training in lifelong learning. In *ICML CL Workshop*.
- Chengwei Qin and Shafiq Joty. 2021. [LFPT5: A unified framework for lifelong few-shot language learning based on prompt tuning of T5](#). *CoRR*, abs/2110.07298.
- Yujia Qin, Jiajie Zhang, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2022. [ELLE: efficient lifelong pre-training for emerging data](#). In *Findings of the Association for Computational Linguistics: ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 2789–2810. Association for Computational Linguistics.
- Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, H. Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan, Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks, Mari-beth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron Huang,

- Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu, Erich Elsen, Siddhant M. Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro, Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch, Jean-Baptiste Lespiau, Maria Tsim-poukelli, Nikolai Grigorev, Doug Fritz, Thibault Sotiaux, Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d'Autume, Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas, Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake A. Hechtman, Laura Weidinger, Iason Gabriel, William S. Isaac, Edward Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol Vinyals, Kareem Ayoub, Jeff Stanway, Lorraine Bennett, Demis Hassabis, Koray Kavukcuoglu, and Geoffrey Irving. 2021. [Scaling language models: Methods, analysis & insights from training gopher](#). *CoRR*, abs/2112.11446.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. 2017. [icarl: Incremental classifier and representation learning](#). In *CVPR*.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 255–269. Association for Computational Linguistics.
- Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. 2017. [Continual learning in generative adversarial nets](#). *CoRR*, abs/1705.08395.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. 2018. [Overcoming catastrophic forgetting with hard attention to the task](#). In *ICML*.
- Yilin Shen, Xiangyu Zeng, and Hongxia Jin. 2019. [A progressive model to enable continual learning for semantic slot filling](#). In *EMNLP-IJCNLP*.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. [Continual learning with deep generative replay](#). In *NIPS*.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, Elton Zheng, Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi, Yuxiong He, Michael Houston, Saurabh Tiwary, and Bryan Catanzaro. 2022. [Using deepspeed and megatron to train megatron-turing NLG 530b, A large-scale generative language model](#). *CoRR*, abs/2201.11990.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. [How to fine-tune bert for text classification?](#) In *China national conference on Chinese computational linguistics*, pages 194–206. Springer.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2020. [Lamol: Language modeling is all you need for life-long language learning](#). In *ICLR*.
- Eleni Triantafillou, Richard S. Zemel, and Raquel Urtasun. 2017. [Few-shot learning through an information retrieval lens](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2255–2265.
- Gido M. van de Ven and Andreas S. Tolias. 2019. [Three scenarios for continual learning](#). <https://arxiv.org/pdf/1904.07734.pdf>.
- Zirui Wang, Sanket Vaibhav Mehta, Barnabás Póczos, and Jaime Carbonell. 2020. [Efficient meta lifelong-learning with limited memory](#). In *EMNLP*.
- Congying Xia, Wenpeng Yin, Yihao Feng, and Philip S. Yu. 2021. [Incremental few-shot text classification with multi-round new classes: Formulation, dataset and system](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1351–1360. Association for Computational Linguistics.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2019. [BERT post-training for review reading comprehension and aspect-based sentiment analysis](#). In *NAACL-HLT*.
- Wenpeng Yin, Jia Li, and Caiming Xiong. 2022. [Contintin: Continual learning from task instructions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3062–3072. Association for Computational Linguistics.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *NIPS*.

A Related Work

Our work is related to *continual learning*, *post-training* and *few-shot learning*.

Continual learning (CL). In general, overcoming CF is a major goal in CL (Chen and Liu, 2018). (1) *Regularization methods* (Kirkpatrick et al., 2016; Seff et al., 2017) add a regularization to ensure minimal changes to weights for previous tasks. (2) *Replay methods* retain (Rebuffi et al., 2017; Lopez-Paz and Ranzato, 2017; Wang et al., 2020; Guo et al., 2022) or generate some data of old tasks (Shin et al., 2017; He and Jaeger, 2018) and use them in learning a new task. (3) *Parameter isolation methods* (Serrà et al., 2018; Fernando et al., 2017) allocate parameters for different tasks and mask them out in learning a new task. Our CPT is based on (3) and uses masks in continual post-training. Recently, CL has drawn attention in NLP. It has been used for slot filling (Shen et al., 2019), language learning (Li et al., 2019), sentence embedding (Liu et al., 2019a), translation (Khayrallah et al., 2018), cross-lingual modeling (Liu et al., 2020b), question answering (Greco et al., 2019) and text classification (Ke et al., 2021a,b; Sun et al., 2020; Huang et al., 2021; Chuang et al., 2020; Mehta et al., 2021; Madotto et al., 2020). However, none of them tries to improve an LM.

Post-training is an effective approach to mitigate the discrepancies between pre-trained domains and the target domain. Researchers have applied post-training to many domains, e.g., reviews (Xu et al., 2019; Sun et al., 2019), news and academic papers (Gururangan et al., 2020b), and shown improved end-task results. However, none of them consider the continual learning paradigm.

Few-shot learning (FL) aims to learn tasks with a few labeled examples. The main issue of FL is over-fitting, due to the scarcity of labeled training data. Existing methods to overcome over-fitting fall in three main families: (i) model-based methods try to reduce the hypothesis space of the few-shot task (Triantafillou et al., 2017; Hu et al., 2018), (ii) data-based methods try to augment additional data to the few-shot set (Benaim and Wolf, 2018; Gao et al., 2020), and (iii) algorithm-based solutions try to improve strategies for searching for the best hypothesis. Recently, a new paradigm using prompts achieves promising results for few-shot language learning as shown in GPT-3 (Brown et al., 2020a), PET (Schick and Schütze, 2021) and LM-BFF (Gao et al., 2021). However, none of

them does few-shot fine-tuning in continual post-training.

Continual few-shot learning. Several researchers have studied this problem recently (Antoniou et al., 2020; Qin and Joty, 2021; Jin et al., 2021; Xia et al., 2021; Yin et al., 2022). It continually learns a sequence of few-shot tasks. However, this is very different from our continual post-training because our continual learning happens in the post-training stage instead of the end-task fine-tuning stage. We only evaluate the proposed CPT system after continual post-training by conducting few-shot learning tasks individually by fine-tuning the post-trained language model (p-LM) in each of the post-trained domains. No continual learning is involved in few-shot learning.

B Illustration of Task Masks

Figure 2 illustrates the CPT architecture and the task mask learning. Note that fine-tuning is for evaluating the domain post-training and should not affect any parameters of post-training. During **continual post-training** (Figure 2 (A)), after training domain/task 1, we obtain its useful neurons indicated by the 1 entries. Before training domain/task 2, those useful neurons for domain 1 are first masked (those previous 1’s entries are turned to 0’s). After training domain 2, two neurons with 1 are used by the domain. When domain t arrives, all used neurons by domains 1 and 2 are masked before training, i.e., their entries are set to 0. After training domain t , we see that domains t and 1 have a shared neuron (the cell with two colors, red and green), which is used by both of domains. After continual post-training, we evaluate CPT by **individual fine-tuning**. During fine-tuning (Figure 2 (B)), we only make use of those neurons that are useful for domain/task id t (red cells) and freeze all other neurons (grey cells).

C Dataset Statistics

Table 3 shows the statistics of the *unlabeled domain datasets* and *end-task classification datasets*. Note that the full AGNews is very large. We use only its author provided training split as our domain-specific datasets as our *unlabeled AGNews* dataset for continual post-training. The remaining testing set is used as the labeled end-task (*AGNews-FT*). The other three corresponding end task datasets are *SemEval-res* (Xu et al., 2019), *ACL-ARC* (Jurgens et al., 2018), and *SCIERC* (Luan et al., 2018).

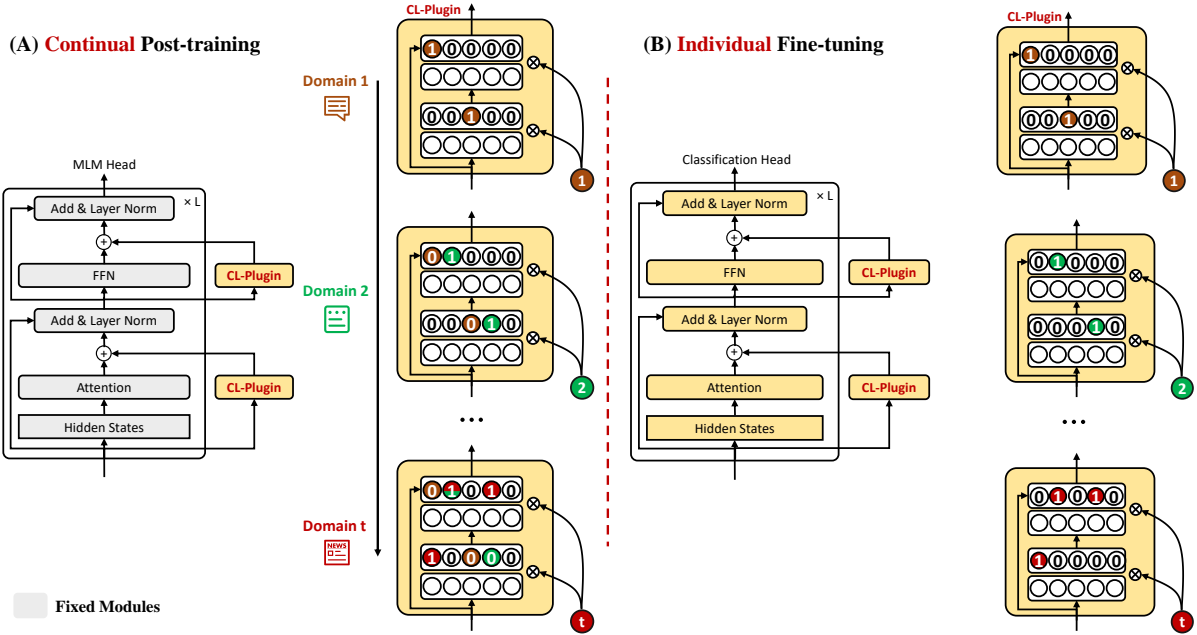


Figure 2: Architecture of CPT, which has two CL-plugins inserted in the transformer layers of RoBERTa in a parallel manner. (A) CPT for *continual post-training*. It uses a masked language model (MLM) head for unsupervised post-training of the plugins only. (B) CPT for *individual fine-tuning*. The performance of CPT is evaluated by the corresponding individual end-task performance of all post-trained tasks using the *final* post-trained model (with different mask). Each CL-plugin module (**to the right of the transformer**) has two fully connected layers and a skip connection. On top of each fully connected layer, there is a mask computed from task ID t with the same size as the fully connected layer.

Unlabeled Domain Datasets			End-Task Classification Datasets				
Dataset	Source	#training	Dataset	Task	#training	#testing	#classes
Yelp Restaurant	Yelp Review	1,132,359	SemEval-res	Aspect Sentiment Classification	32	1,120	3
AI	AI Papers	707,368	SCIERC	Relation Classification	56	2,388	7
ACL	ACL Papers	1,208,449	ACL-ARC	Citation Intent Classification	48	421	6
AGNews	News Article	73,750	AGNews-FT	News Classification	32	7,568	4

Table 3: Statistics for unlabeled domain datasets and end task supervised classification datasets.

D Details of the CL baselines

Non-Continual Learning Baselines: Each of these baselines builds a separate model for each task independently. It thus has no CF.

(1,2) **RoBERTa, Adapter** (Liu et al., 2019b; Houlsby et al., 2019) use the original RoBERTa/Adapter for the end-task fine-tuning without any post-training. These are the only two without any post-training. All the following baselines use the masked language model loss (MLM) for post-training.

(3) **RoBERTa-ONE** is the existing post-training method in (Gururangan et al., 2020b). To our knowledge, the existing post-training systems are all based on the MLM loss.

(4) **Adapter-ONE** (Madotto et al., 2020; Houlsby et al., 2019) adds small adapter layers

between layers of Transformer for post-training. We follow the adapter design in (Madotto et al., 2020; Houlsby et al., 2019). An adapter is simply two fully connected layers. During post-training, the Transformer is fixed, only the added adapters are trainable. The bottleneck size (adapter size) is set to 128. During end-task fine-tuning, both RoBERTa and the adapters are trainable to ensure fair comparison.

(5) **Prompt-ONE** (Lester et al., 2021) adds a sequence of real vector tokens (called virtual tokens or prompt tokens) to the end of the original input sequence. In post-training, RoBERTa (the LM) is fixed and only the prompt tokens are trained. In end-task fine-tuning, both LM and the trained prompt are trainable. We initialize 100 tokens and set the learning rate of the prompt token to 0.3 in

Category	Domain Model	Restaurant		AI		ACL		AGNews		Average	
		MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc
Non-CL	RoBERTa	±0.0456	±0.0274	±0.0208	±0.0233	±0.0391	±0.0338	±0.0121	±0.014	±0.0066	±0.0062
	Adapter	±0.0214	±0.0223	±0.0111	±0.0102	±0.0375	±0.0386	±0.0221	±0.0224	±0.0155	±0.0142
	RoBERTa-ONE	±0.0095	±0.0087	±0.0364	±0.0358	±0.0382	±0.0432	±0.0169	±0.0162	±0.0197	±0.0187
	Adapter-ONE	±0.0292	±0.0223	±0.0207	±0.0222	±0.0076	±0.0063	±0.0141	±0.0157	±0.0074	±0.0054
	Prompt-ONE	±0.0427	±0.0991	±0.0297	±0.0254	±0.0386	±0.0325	±0.0115	±0.0100	±0.0151	±0.0292
	DEMIX	±0.0329	±0.0293	±0.0259	±0.0283	±0.0297	±0.0367	±0.0336	±0.0309	±0.0152	±0.0165
CL	RoBERTa-NCL	±0.0374	±0.0238	±0.0156	±0.0158	±0.0293	±0.0349	±0.0218	±0.0154	±0.0130	±0.0160
	Adapter-NCL	±0.0250	±0.0194	±0.0232	±0.0184	±0.0183	±0.0264	±0.0136	±0.0151	±0.0095	±0.0137
	HAT	±0.0264	±0.012	±0.0236	±0.0251	±0.0294	±0.0287	±0.0106	±0.009	±0.0078	±0.0112
	BCL	±0.0255	±0.0124	±0.0121	±0.0105	±0.0182	±0.0126	±0.0100	±0.0069	±0.0094	±0.0032
	KD	±0.0642	±0.0435	±0.0295	±0.0233	±0.0271	±0.0267	±0.0160	±0.0133	±0.0117	±0.0109
	EWC	±0.0324	±0.0259	±0.0281	±0.0189	±0.0177	±0.0196	±0.0041	±0.0096	±0.0079	±0.0062
	DER++	±0.0250	±0.0183	±0.0231	±0.0319	±0.0116	±0.0163	±0.0196	±0.0178	±0.0126	±0.0128
	CPT	±0.0264	±0.0120	±0.0236	±0.0251	±0.0294	±0.0287	±0.0106	±0.0090	±0.0078	±0.0112

Table 4: Standard deviations of the corresponding metrics of the proposed CPT system and the baselines.

Model	Final Performance	
	MF1	Acc
CPT (Sequential Adapter)	±0.0347	±0.0350
CPT (w/o butterfly)	±0.0102	±0.0079
CPT (w/o masking)	±0.0095	±0.0137
CPT	±0.0078	±0.0112

Table 5: Standard deviations of the corresponding metrics of the proposed CPT system and the ablations.

post-training, following the setting in (Lester et al., 2021).

(6) **DEMIX** (Gururangan et al., 2021) is a recent model to adapt a pre-trained LM with new domains. It adds a new adapter once a new domain arrives (network expansion is needed) and initializes the new adapter with the parameters of the previous trained adapter nearest to the new domain data. They use the perplexity on held-out samples to choose the most probable adapter.

Continual Learning (CL) Baselines.

(7) **RoBERTa-NCL (Naive continual learning)** is a naive extension of (Gururangan et al., 2020b), which continually/incrementally post-trains the LM to learn all domains using the MLM loss with no mechanism to deal with forgetting or CF.

(8) **Adapter-NCL** (Houlsby et al., 2019) is similar to the Adapter based system. The only difference is that the same set of adapters is shared across all domains, rather than using a new adapter for each new domain.

(9) **Hard attention to overcome forgetting (HAT)** is derived from HAT (Serrà et al., 2018), the state-of-the-art parameter-isolation based method with almost no forgetting. However, HAT suffers from forgetting in continual post-training due to the catastrophic butterfly effect.

(10) **BCL** (Ke et al., 2021c) is a continual learn-

ing model that can avoid forgetting and encourage knowledge transfer. It is similar to Adapter-NCL. The difference is that its adapters consist of two modules, one is a capsule network (a new capsule is added once a new domain arrives) to encourage transfer, and the other is similar to HAT to avoid forgetting. Similar to HAT, task/domain information is needed in end-task fine-tuning. We replace the backbone network from BERT with RoBERTa for fair comparison.

(11) **Knowledge distillation (KD)** (Hinton et al., 2015) minimizes the representational deviation between the learned representation and the new representation in post-training. We compute the KL divergence between the representations (the output before the masked language model prediction head) of each token of the previous post-trained LM and current LM as the distillation loss.

(12) **EWC** (Buzzega et al., 2020) is a popular regularization-based continual learning method that adopts elastic weights consolidation to add L_2 regularization to penalize parameter changes.

(13) **DER++** (Buzzega et al., 2020) is a recent replay method using distillation to regularize the new task training. We store 16.4K tokens for each learned domain as the memory, which is the largest memory we can use for the system to run.

E Results for Different Domain Orders

Table 1 in the main paper reported the results for the order Restaurant \rightarrow AI \rightarrow ACL \rightarrow AGnews. We now look at how the order affects the results. Table 6 shows baselines and CPT’s results of 4 different orders. Note that the results for the Non-CL baselines are the same across different orders (and the same as those in Table 1) because they

Category	Order Metric Model	AI→ACL→Restaurant→AGNews				Restaurant→AI→AGNews→ACL				AI→ACL→AGNews→Restaurant				AGNews→ACL→Restaurant→AI				Average	
		Performance		Forget R.		Performance		Forget R.		Performance		Forget R.		Performance		Forget R.		Performance	Forget R.
		MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc	MF1	Acc
Non-CL	RoBERTa	43.72	50.94	—	—	43.72	50.94	—	—	43.72	50.94	—	—	43.72	50.94	—	—	43.72	50.94
	Adapter	39.65	46.48	—	—	39.65	46.48	—	—	39.65	46.48	—	—	39.65	46.48	—	—	39.65	46.48
	RoBERTa-ONE	44.78	51.92	—	—	44.78	51.92	—	—	44.78	51.92	—	—	44.78	51.92	—	—	44.78	51.92
	Adapter-ONE	45.31	51.68	—	—	45.31	51.68	—	—	45.31	51.68	—	—	45.31	51.68	—	—	45.31	51.68
	Prompt-ONE	34.68	43.42	—	—	34.68	43.42	—	—	34.68	43.42	—	—	34.68	43.42	—	—	34.68	43.42
	DEMIX	45.41	51.57	—	—	45.41	51.57	—	—	45.41	51.57	—	—	45.41	51.57	—	—	45.41	51.57
CL	RoBERTa-NCL	42.62	49.95	2.45	1.79	42.22	49.52	3.10	2.33	42.88	50.11	0.29	0.18	44.33	51.51	1.76	1.21	43.01	50.28
	Adapter-NCL	44.71	51.67	1.71	1.08	44.61	51.07	1.14	1.23	44.91	51.57	1.41	1.23	45.52	52.15	0.72	0.44	44.94	51.62
	HAT	45.10	51.50	1.66	1.19	43.29	49.96	2.76	2.09	46.06	52.07	0.50	0.21	44.94	51.45	0.86	0.25	44.85	51.25
	BCL	43.97	50.74	2.20	1.50	45.30	51.54	0.36	-0.14	45.28	51.79	0.36	0.11	45.59	51.61	0.08	0.11	45.04	51.42
	KD	42.09	50.22	0.57	0.08	45.18	52.68	1.22	0.57	42.63	50.45	-0.31	-0.56	42.93	50.70	1.10	0.32	43.21	51.01
	EWC	43.97	50.74	0.16	0.03	43.65	50.29	-0.29	-0.20	45.52	51.36	0.17	0.15	43.42	49.85	0.12	0.10	44.14	50.56
	DER++	44.56	50.13	2.95	2.31	44.02	49.99	1.24	1.12	43.98	50.23	1.44	1.27	44.32	50.13	1.32	1.09	44.22	50.12
	CPT	46.49	52.47	0.00	0.00	45.71	51.71	0.00	0.00	46.15	51.93	0.00	0.00	45.89	51.86	0.00	0.00	46.06	51.99

Table 6: CPT performance averaged over all domains after the final post-trained with different orders (averaged over 5 random seeds) and the average of these orders.

are not effected by orders. We can see CPT is always better than other baselines, and achieve 0 forgetting rate, demonstrating the effectiveness of CPT. We also note that some baselines in some sequence has negative forgetting rate, indicating they have some backward transfer (new domain learning helps learned domains). However, their final results are much worse than CPT’s.

F Standard Deviations

Table 4 reports the standard deviations of the corresponding results in Table 1 (in the main paper) of CPT and the considered baselines over 5 runs with random seeds. We can see the results of CPT are stable. Some baselines (e.g., RoBERTa, RoBERTa-ONE) can have quite large standard deviations.

Table 5 reports the standard deviations of the corresponding results in Table 2 (in the main paper) of CPT and the considered baselines over 5 runs with random seeds. We can see the results of sequential adapters has a high variance while CPT and other variants are stable.