

Learning Label Modular Prompts for Text Classification in the Wild

Hailin Chen^{♣♣}, Amrita Saha[♣], Shafiq Joty^{♣♣}, Steven C.H. HOI[♣]

[♣] Nanyang Technological University, Singapore

[♣] Salesforce Research

{hailin001, srjoty}@ntu.edu.sg

{amrita.saha, shoi}@salesforce.com

Abstract

Machine learning models usually assume i.i.d data during training and testing, but data and tasks in real world often change over time. To emulate the transient nature of real world, we propose a challenging but practical task: text classification *in-the-wild*, which introduces different non-stationary training/testing stages. Decomposing a complex task into modular components can enable robust generalisation under such non-stationary environment. However, current modular approaches in NLP do not take advantage of recent advances in parameter efficient tuning of pretrained language models. To close this gap, we propose MODULARPROMPT, a label-modular prompt tuning framework for text classification tasks. In MODULARPROMPT, the input prompt consists of a sequence of soft *label* prompts, each encoding modular knowledge related to the corresponding class label. In two of most formidable settings, MODULARPROMPT outperforms relevant baselines by a large margin demonstrating strong generalisation ability. We also conduct comprehensive analysis to validate whether the learned prompts satisfy properties of a modular representation.¹

1 Introduction

While NLP research has received a significant boost in performance by employing large-scale pretrained language models (PLMs), finetuning an entire dedicated model for each task is not always practical or even feasible, especially as model size continues to grow. To alleviate this, recently there has been increased interest in parameter-efficient methods such as Adapter (Houlsby et al., 2019) and Prompt-based tuning methods (Lester et al., 2021; Qin and Eisner, 2021; Liu et al., 2021; Li and Liang, 2021). When training on a downstream task, these methods keep the PLM frozen and only

¹Our Code is available at <https://github.com/salesforce/ModularPrompt>



Figure 1: Example of MODULARPROMPT for stage-fused NER. **Top 2 blocks**: training stages covering disjoint sets of entity types (e.g., *event*, *person*). **Bottom block**: fused test stage covering entity types of *person* and *organization* from two training stages. Coloured boxes denote label prompts. Underlining and italic blue denotes named entity and its type.

update a small set of parameters that are added to the network. Among these methods, PROMPTTUNING (Lester et al., 2021) proposes *tunable prompts* – a sequence of learnable soft tokens, and show impressive results, even competitive to full-model finetuning with a large PLM.

While these prompt-based methods have proven quite effective on popular benchmarks, these standard tasks typically assume only independently and identically distributed (i.i.d) data during training and testing. However, practical cognitive tasks in the real world are usually more complex involving changing contexts or non-stationary environments. Talking about what is next for NLP, Kathleen McKown in a recent interview (source) said: “Most models are static. But the world changes every minute, every second. Dealing with a dynamic world is a new area that’s up and coming.”

Our work in this paper particularly concerns text classification settings where a model is trained on

sequence of tasks and evaluated on an arbitrary subset of seen labels of interest. We formalize this as a novel **text classification *in-the-wild*** task (defined in §3), which emulates the transient learning environment of real world, e.g., for a service requiring classification, the label set might gradually change over time to include new labels or remove obsolete ones. Such scenarios typically result in a sequence of non-stationary low-resource training and evaluations over different label sets (e.g., train on {*chemistry*, *physics*} and {*basketball*, *football*} in succession and then test on {*physics*, *football*}).

This requires handling non-stationary data distribution which humans are quite adept at, partly because we can decompose a complex task in a modular fashion (Berwick et al., 2013). For example, when learning to classify objects, we acquire modular knowledge exclusive to each class. This allows us to robustly classify irrespective of any label space manipulations such as label omission or learning over new label spaces. This notion of modularity at the level of each class label is what we call *label modularity* and is a desirable quality for NLP models to generalize to practical non-stationary classification settings.

Contemporary modular model designs for complex NLP tasks typically use a routing network or programmer (Cases et al., 2019; Rosenbaum et al., 2019; Khot et al., 2021; Corona et al., 2021; Jiang and Bansal, 2019; Liu et al., 2019; Hu et al., 2018; Gupta et al., 2020; Andreas et al., 2016) which learns a meaningful decomposition of the task into sub-tasks and executes it by applying a chain of specialised modules designed for each sub-task. For classification tasks, this can entail learning a specialized module for each label in the target label-space. While there has been limited research on utilizing modular architectures for PLMs (Andreas et al., 2016; Chen et al., 2020), the main research gap that we explore in this work is that of a modular design of parameter efficient tuning of large PLMs, in particular, PROMPTTUNING.

Although PROMPTTUNING can be considered modular at task level in that it learns soft-prompts for each task to support multitasking, it is not able to learn modular decomposition within a particular task. For non-modular designs like PROMPTTUNING or model-finetuning, text classification *in-the-wild* is challenging to handle, as it requires combining partial information from different label spaces. In contrast, a label-modular approach should learn ex-

clusive knowledge for each label and generalise to any subset of the label set. We thus postulate two main objectives of a label-modular model:

Objective 1. Separable Label Representation: *Each class label should have its own representation which compactly encodes the information from the data belonging to that label.*

Objective 2. Prediction over Controllable Label Space: *Models should perform robustly over any subset of the learnt label space during inference.*

To meet these objectives, we propose a modular design of Prompt Tuning – **Label-modular Prompt Tuning** (MODULARPROMPT). It decomposes the prompt sequence into label-modular components called *label prompts*, each encoding specific knowledge corresponding to a class label. Thus in each forward pass, we can select desired label prompts to construct the input prompt, based on the target label-set. To ensure that the learned knowledge is encoded in a modular fashion during training, we introduce a novel subset-invariant loss over dynamic label-sets.

To evaluate generalizability of MODULARPROMPT we construct some practical scenarios of text classification *in-the-wild*. We train in multiple stages over non-overlapping label spaces and evaluate the model on label-sets that (i) correspond to each training stage (stage-specific), (ii) is accumulated over all learned labels (stage-agnostic), and (iii) is comprised of labels across multiple training stages (stage-fused). We show an example of MODULARPROMPT on stage-fused NER setting in Figure 1.

The stage-agnostic and stage-fused settings are the most challenging scenarios for typical fine-tuned or prompt-tuned models, and specifically on those settings we find that MODULARPROMPT outperforms all relevant baselines by a significant margin. This alludes towards its ability to learn robust prompt representations that is generalizable to different non-stationary learning environments.

We further empirically justify that MODULARPROMPT indeed showcases modular properties by analyzing its behavior when either the ground truth or other random labels are removed from the input or the order of label prompts is permuted.

2 Related Work

We now review methods from the literature that are relevant to ours from different perspectives. First, many parameter efficient tuning methods have been proposed such as Adapter (Houlsby et al., 2019),

PROMPTTUNING (Lester et al., 2021), PREFIXTUNING (Li and Liang, 2021), BitFit (Zaken et al., 2022), LoRA (Hu et al., 2021) and COMPACTER (Mahabadi et al., 2021a). They either finetune a parameter subset or introduce new layers or embeddings. Transfer learning over prompts (Vu et al., 2022) and Adapters (Pfeiffer et al., 2021) and multi-tasking over the latter (Mahabadi et al., 2021b) have also been explored, but they are not comparable to ours due to difference in task/problem settings.

Second, continual learning (CL) methods are relevant as they also have sequential training stages. Architecture based CL methods adjust model architecture for each task (Chen et al., 2016; Rusu et al., 2016; Mallya et al., 2018), but require task identities for inference. Regularization based CL methods restrain updating parameters critical to previous tasks (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018). Memory based CL methods retain key examples from prior tasks (Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2019; de Masson d’Autume et al., 2019; Zhu et al., 2022) while Memory generator models learn to generate and use pseudo-data from prior tasks (Sun et al., 2020; Qin and Joty, 2021). These methods are not comparable to ours as we do not use any memory or pseudo memory in our sequential training.

Third, modular networks have been shown to perform well on out-of-domain data (Kirsch et al., 2018; Ruder et al., 2019; Alet et al., 2018) and mitigate forgetting in continual learning (Ostapenko et al., 2021). Apart from routing network approaches introduced in §1, mixture-of-experts (MoE) selects a soft subset of modules based on model input (Shazeer et al., 2017; Fedus et al., 2021; Goyal et al., 2021). Similar to ours, Kudugunta et al. (2021); Rajendran et al. (2017); Ponti et al. (2021); Ostapenko et al. (2021) consider task level routing and support new tasks by combining learned modules. By contrast, MODULARPROMPT utilises parameter efficient finetuning of large PLMs and support low resource settings.

3 Methodology

In this section, we first formally define the problem and subsequently present our MODULARPROMPT model, introduce subset invariant loss and explain the framework under text classification *in-the-wild*.

3.1 Problem Definition

Single stage text classification Assume a single text classification domain (or dataset) \mathcal{D} . Let $(X, Y) \sim \mathcal{D}$ be a sample, where $X = \{x_t\}_{t=1}^L$ represents a text input sequence of length L and $Y = \{y_t\}_{t=1}^M$ represents the corresponding classification label name of length M (in tokens). Let Ω denote the set of all possible class labels of interest, for which we have $\forall (X, Y) \sim \mathcal{D}, \text{cls}(Y) \subseteq \Omega$. Note that $\text{cls}(Y)$ is a mapping which returns the class label(s) in Y . In case of single class classification, $\text{cls}(Y)$ returns $\{Y\}$. In case of sequence labelling which is token-level classification, $\text{cls}(Y)$ returns the set of all unique target tags in Y .

Text classification in-the-wild Assume a sequence of n text classification stages with the corresponding training datasets $\mathbb{D}^{tr} = \{\mathcal{D}_1^{tr}, \dots, \mathcal{D}_n^{tr}\}$. Each stage represents a different task in temporal dimension, with $(X_k, Y_k) \sim \mathcal{D}_k^{tr}$ denoting a sample at the k -th training stage and Ω_k denoting the set of all possible class labels for \mathcal{D}_k^{tr} . Similarly, the testing could consist of m such datasets $\mathbb{D}^{ts} = (\mathcal{D}_1^{ts}, \dots, \mathcal{D}_m^{ts})$ with Ω_j^{ts} denoting the set of possible class labels for \mathcal{D}_j^{ts} . For classification *in-the-wild*, we examine three challenging yet very practical settings. First, when $m = 1$ and $\Omega_1^{ts} = \cup_{k=1}^n \{\Omega_k^{tr}\}$, we have one test dataset covering all seen labels. We refer it as **stage-agnostic** testing as the test label can come from any of the training stages (or tasks). Second, when $m = n$ and $\Omega_j^{ts} = \Omega_j^{tr}, \forall j = \{1, \dots, n\}$, we have one test dataset corresponding to each training stage with the same label set. We denote this setting as **stage-specific** testing as each test set evaluates the model’s performance on a particular task in which it was trained. Finally, a more challenging setting where $m > 1$ and $\Omega_j^{ts} \notin \{\Omega_1^{tr}, \dots, \Omega_n^{tr}\}, \forall j = \{1, \dots, m\}$, rather $\Omega_j^{ts} \in \mathcal{P}(\cup_{k=1}^n \{\Omega_k^{tr}\}) - \cup_{i=1}^n \{\mathcal{P}(\Omega_i^{tr})\}$, where $\mathcal{P}(S)$ denotes the power-set of a given set S . That is, the label set of a test stage does not correspond to any one training stage, but is composed of partial label sets from multiple training stages (or tasks). We refer it as **stage-fused** testing. Note that the stage-agnostic and stage-specific scenarios are closely related to *continual learning* (Thrun, 1996), though the latter considers access to task-id instead of intra-task information (i.e., task label set).

3.2 Soft Prompt Tuning

Let $X = \{x_1, \dots, x_L\}$ be an input text sequence, where x_t is the t -th token, and M be a pretrained

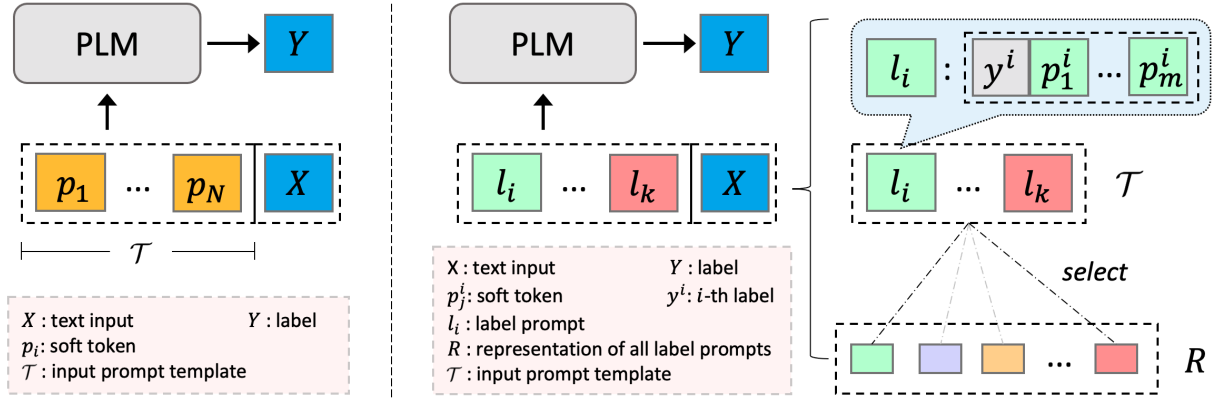


Figure 2: **Left:** PROMPTTUNING, where consecutive soft tokens are concatenated with the input to a frozen PLM. **Right:** Our proposed MODULARPROMPT, where multiple label prompts are concatenated together with the input to a frozen PLM. Each label prompt consists of a label name and consecutive sequence of soft tokens.

language model. The input text is mapped to a sequence of embeddings $H = \{h_1, \dots, h_L\}$ with $h_t \in \mathbb{R}^d$. A soft prompt is a sequence of N tunable tokens $\mathcal{T} = \{p_1, \dots, p_N\}$ with $p_i \in \mathbb{R}^d$, that is concatenated with the text embedding as the final input to M : $\bar{H} = \{\mathcal{T} \oplus H\} = \{p_1, \dots, p_N, h_1, \dots, h_L\}$. The model prediction is defined as $P(Y|\bar{H}; M) = P(Y|\mathcal{T}, X; M)$. During training, M is kept frozen and only \mathcal{T} is updated.

3.3 Label Modular Prompt Model

We visualise the architecture of PROMPTTUNING (Lester et al., 2021) and our proposed MODULARPROMPT in Figure 2. In contrast to PROMPTTUNING, MODULARPROMPT’s prompt consists of a sequence of label prompts, where each label prompt contains the corresponding label name and a sequence of tunable soft tokens similar to soft prompt. Formally, we denote $l_k = e_k \oplus \{p_1^k, \dots, p_m^k\}$ as label prompt for label k , where e_k is the embedding of label k ’s text or sequence of token-embeddings for multi-token labels, \oplus denotes concatenation and m is the number of tunable tokens per label prompt. The final input prompt is $\mathcal{T} = \oplus_{k \in S} l_k$ with S being the set of labels of interest.

Prompt formulation The key mechanism of MODULARPROMPT is prompt formulation $\{R, S\} \rightarrow \mathcal{T}$, where R denotes the learned representation space of all labels prompts. In PROMPTTUNING, variables S and R do not exist and the model training tunes \mathcal{T} directly. In MODULARPROMPT, given S as a set of class labels of interest, we select the corresponding label prompts representation from R and concatenate these to form the final input prompt \mathcal{T} . The training loss is back-propagated through $Y \rightarrow \mathcal{T} \rightarrow R$ to learn the soft label prompts.

Subset invariant loss The prompt formulation $\{R, S\} \rightarrow \mathcal{T}$ aims to achieve Objective 2: prediction over controllable label space (§1). In single domain setting, Ω^{tr} is the set of all possible class labels during training as defined by Section 3.1. However fixing S to a constant Ω^{tr} throughout training will make the model susceptible to data discrepancy between train and inference as $\Omega^{ts} \neq \Omega^{tr}$. Thus to ensure Objective 2, we propose to vary S during training. We first uniformly sample the size of S , $|S|$ from $\{1, \dots, (|\Omega^{tr}| - 1)\}$ and then randomly choose $|S|$ labels from Ω^{tr} to construct S . Such sub-sampling of Ω^{tr} encourages a fair exploration of different lengths of prompt sequences as input during training, thus enabling representations to be robust to a dynamic Ω^{ts} at inference. For each training instance, with probability p we fix $S = \Omega^{tr}$ and vary S as above with $(1 - p)$ chance. We refer such sampling process as $S \sim \hat{S}$. The subset invariant loss is then defined as:

$$\mathcal{L}_R^{\text{inv}}(\mathcal{D}) = \mathbb{E}_{\substack{(X, Y) \sim \mathcal{D} \\ S \sim \hat{S}}} [-\mathbb{1}_{\text{cls}(Y) \subseteq S} \log P(Y|X, R, S; M)] \quad (1)$$

where $\mathbb{1}$ is the Indicator function; $\mathbb{1}_{\text{cls}(Y) \subseteq S} = 1$ if $\text{cls}(Y) \subseteq S$, otherwise 0. According to Objective 1, we expect our model to make predictions grounded by the relevant label prompts. When S does not contain ground truth class label(s) in Y , the model should not be able to predict Y as output. Thus we set the loss to be zero when $\text{cls}(Y) \not\subseteq S$ to avoid encouraging ungrounded predictions.

3.4 Classification *in-the-wild* with MODULARPROMPT

So far, we have introduced MODULARPROMPT’s functionality under a single domain classification setting. To verify our Objective 2, we wish to examine it under text classification *in-the-wild* defined in

Algorithm 1 Text Classification *in-the-wild* with MODULARPROMPT.

Require: Training datasets $\mathbb{D}^{tr} = \{\mathcal{D}_1^{tr}, \dots, \mathcal{D}_n^{tr}\}$, testing datasets $\mathbb{D}^{ts} = (\mathcal{D}_1^{ts}, \dots, \mathcal{D}_m^{ts})$, pretrained language model M

- 1: **for** \mathcal{D}_i^{tr} in \mathbb{D}^{tr} **do** ▷ Training
- 2: $\Omega_i^{tr} \leftarrow \Omega_i^{tr}$
- 3: INITIALIZE $R_{\Omega_i^{tr}}$ from $R_{\Omega_{<i}^{tr}}$ ▷ Transfer
- 4: **for** $step = 1, \dots, iter$ **do**
- 5: UPDATE $R_{\Omega_i^{tr}}$ with SGD on $\mathcal{L}_R^{inv}(\mathcal{D}_i^{tr})$
- 6: **end for**
- 7: **end for**
- 8: **for** \mathcal{D}_i^{ts} in \mathbb{D}^{ts} **do** ▷ Testing
- 9: EVALUATE $P(\cdot|X, S = \Omega_i^{ts}, R; M)$ on $X \sim \mathcal{D}_i^{ts}$
- 10: **end for**

§3.1. Given training datasets $\mathbb{D}^{tr} = \{\mathcal{D}_1^{tr}, \dots, \mathcal{D}_n^{tr}\}$, the model is trained on each dataset \mathcal{D}_i^{tr} sequentially, and then evaluated on three classification *in-the-wild* testing settings. The pseudocode of MODULARPROMPT under the continual learning setting is given in Algorithm 1. Note that $R_{\Omega_i^{tr}}$ in step 3 denotes label prompt representation of labels in Ω_i^{tr} , i.e., $R_{\Omega_i^{tr}} := \{\mathbf{l}_k \in R | k \in \Omega_i^{tr}\}$ and $R_{\Omega_{<i}^{tr}}$ is similarly defined as $R_{\Omega_{<i}^{tr}} := \{\mathbf{l}_m \in R | m \in \bigcup_{t < i} \Omega_t^{tr}\}$.

Label prompt transfer In step 3, for learning the label prompt representation $R_{\Omega_i^{tr}}$ at any training stage i , we first aim to transfer the label-modular knowledge, $R_{\Omega_{<i}^{tr}}$ learned over the previous training stages through prompt initialization. This is a unique learning characteristic that is facilitated by our label-modular architecture and allows the model to exploit semantic relatedness between labels across training stages when initializing the label prompt representation. Intuitively, if ‘bistro’ $\in \Omega_{<i}^{tr}$ and ‘restaurant’ $\in \Omega_i^{tr}$, then initializing the label prompt representation of ‘restaurant’ with the knowledge encoded in the learned label prompt representation of ‘bistro’ should be helpful to the model. To compute the similarity between labels \mathbf{l}_j and \mathbf{l}_k with $j \in \Omega_i^{tr}$ and $k \in \Omega_{<i}^{tr}$, we use per-token average cosine similarity $sim(e_j, e_k)$ based on the embeddings of the label texts. For each label $j \in \Omega_i^{tr}$, we select the top- K most similar labels $\Omega_{top-K(j)}^{tr} \subset \Omega_{<i}^{tr}$. We then initialize \mathbf{l}_j by averaging the top- K similar label prompt representations, weighted by their normalized similarity score: $\mathbf{l}_j \leftarrow \sum_{k \in \Omega_{top-K(j)}^{tr}} \alpha_k \mathbf{l}_k$, where $\alpha_k = sim(e_j, e_k) / \sum_{m \in \Omega_{top-K(j)}^{tr}} sim(e_j, e_m)$. This method is similar in spirit to (Vu et al., 2022), which shows good transfer for task level prompts with training overheads, while we transfer at a finer-

grained level over label prompts with no overheads.

4 Experiments

In this section, we first introduce datasets used and data construction process (§4.1) followed by relevant baselines (§4.2), evaluation methods (§4.3) and implementation details (§4.4). Through our experiments, we target three research questions:

1. Can MODULARPROMPT consolidate knowledge over multi-stage training? → answered in §4.5 with stage-agnostic setting
2. Can MODULARPROMPT adapt to dynamic label space at inference? → answered in §4.6 with stage-fused setting
3. How competitive is MODULARPROMPT in stage-specific setting? → answered in §4.7

Additionally, we perform ablations (§4.8) and quantitative and qualitative analysis (§4.9-§4.10) to verify modular properties of MODULARPROMPT.

4.1 Tasks and Datasets

We conduct experiments on three types of NLP tasks: News Domain Classification on HuffpostNews (Misra, 2018), Name Entity Recognition (NER) on fewNERD (Ding et al., 2021) and Relation Extraction (RE) on FewRel (Han et al., 2018). We formulate all tasks as a text-to-text problem, as defined in §3.1. For News Domain Classification and NER, we construct target text following Qin and Joty (2021). For RE, we concatenate the original text and entities with a separator ‘|’ as the input sequence, and use the relation type as the target. (Example data can be found at Appendix A)

For HuffpostNews, we subsample 100 shots per class for training and validation and split it into 5 stages of disjoint labels. For FewNERD and FewRel, we subsample 50 shots for training and validation and split into 4 and 5 stages, respectively. For testing, we subsample 200, 50, and 50 shots per class for HuffpostNews, FewNERD and FewRel, respectively. The total number of labels for {HuffpostNews, FewNERD, FewRel} is {41, 64, 80} respectively, and resulting label size per stage is {8-9, 16, 16} respectively.

For stage-specific testing, we follow the stages defined for training and construct a corresponding test data for each stage. For stage-agnostic testing, we combine stage-specific test data for current stage and all previously seen stages to construct the test data. For stage-fused testing, we construct label-sets for each fused stage such that

it is not a subset of any single prior training stage, but rather contains labels from ‘all’ prior training stages. We construct {5,4,5} fused stages for {HuffpostNews, FewNERD, FewRel}. We conduct 5 randomised trials with different data sampling and experiment seed for all of the above settings.

4.2 Baselines

We use T5-large (Raffel et al., 2020) as the backbone PLM for all methods, and consider the following baselines to compare with our MODULARPROMPT:

- MODEL TUNING (Finetune), which tunes all parameters of the backbone PLM.
- (i) PROMPT TUNING (PT) from §3.2, (ii) PT_{CL} - An extension of PT to continual learning (CL) setting, which trains separate PT models for each stage and concatenates the learned soft-prompts during inference, based on the test label-set.
- Adapter, a parameter efficient tuning alternative introduced in (Houlsby et al., 2019), which inserts light adapter layers into the backbone PLM and only tune them.

As text classification *in-the-wild* overlaps with continual learning, we also compare with versions of the above baselines that use architecture-agnostic methods and settings relevant to the latter.

- Online regularization based methods: (i) A scalable online version of EWC (Kirkpatrick et al., 2017) proposed in (Schwarz et al., 2018), and (ii) Online MAS (Aljundi et al., 2018). These methods measure each parameter’s importance to previous tasks by fisher information, and restrict updating previously important parameters when learning a new task, to mitigate catastrophic forgetting.
- Multitask model, which involves training on all stages simultaneously, not sequentially. This is in fact an oracle method for stage-agnostic testing and can be considered as an upper bound of memory-based methods in continual learning.

4.3 Evaluation Methods

For all the three NLP tasks, we consider an *exact match* as a correct prediction and report accuracy for News Classification and RE, and compute F1-score over the BIO format for the NER task. By default, we do not apply any other post-processing or verbalizer, though these are orthogonal methods that can be separately used to enhance any of the discussed models. In the stage-fused setting, we

apply constrained decoding similar to (Cao et al., 2021) to selected baselines, marked by special indicator * (e.g., Finetune*_{MAS}). For MODULARPROMPT, we use all seen label prompts for stage-agnostic testing and specific set of label prompts for stage-specific and stage-fused testing. Since other baselines do not have label-level modularity, for stage-agnostic and stage-fused testing, we use the checkpoint after the final stage and for stage-specific testing we take their checkpoints after each training stage. We show average performance in the main paper and relegate detailed results to Appendix C.

4.4 Implementation Details

We set the learning rate to 0.5 for PROMPT TUNING and MODULARPROMPT and 5e-5 for MODEL TUNING and Adapter, using Adafactor (Shazeer and Stern, 2018) optimizer. We adopt implementation of Adapter from OpenDelta (Hu, 2022) and use the default bottleneck dimension of 24. For online EWC and MAS, we report best results obtained over different regularization constant. For all methods, we set maximum training epochs to 256 for HuffpostNews and FewNERD, and to 512 for FewRel. For MODULARPROMPT, the number of soft tokens per label prompt is set to 10, the selection probability p is set to 50% and number of label transfer candidates K in §3.4 is set to 3.

4.5 Results on Stage-agnostic Setting

In Table 1, we show the stage-agnostic testing results. We observe that across all three tasks, MODULARPROMPT significantly outperforms all other baselines by a large margin. This empirically justifies that MODULARPROMPT is indeed able to dynamically combine the label-specific knowledge learned across different training stages in order to infer over the *unseen* combined label-space. Amongst the baselines, MODEL TUNING performs relatively better, while the limited trainable parameters make the parameter efficient models more susceptible to catastrophic forgetting. For CL methods, MAS improves MODEL TUNING and PROMPT TUNING by 4% and 8% on average respectively, but fails on Adapter. EWC is less effective in addressing forgetting across all baselines.

Also note that the PT_{CL} extension is able to improve by 10-20% over vanilla PT. This shows that soft prompts, behaving like language tokens, have a compositional nature and can be concatenated to support multi-tasking. MODULARPROMPT, in addition to exploiting this implicit language prior, also

Methods	Size	News	NER	RE
Finetune _{multitask}	770M	60.6 \pm 5.0	64.5 \pm 0.2	87.9 \pm 0.6
Finetune	770M	23.0 \pm 0.6	25.0 \pm 2.7	40.8 \pm 2.2
Finetune _{EWC}	770M	22.5 \pm 1.2	26.6 \pm 2.8	45.4 \pm 3.2
Finetune _{MAS}	770M	29.2 \pm 2.9	25.4 \pm 2.9	46.2 \pm 1.8
Adapter	4.8M	19.3 \pm 0.3	21.2 \pm 2.0	28.6 \pm 1.9
Adapter _{EWC}	4.8M	18.6 \pm 1.0	20.5 \pm 1.3	27.8 \pm 1.2
Adapter _{MAS}	4.8M	20.1 \pm 1.5	21.1 \pm 2.0	28.3 \pm 0.5
PT	0.13M	16.1 \pm 0.3	18.8 \pm 0.1	19.6 \pm 0.1
PT _{EWC}	0.13M	16.5 \pm 0.5	19.0 \pm 0.5	20.1 \pm 0.8
PT _{MAS}	0.13M	23.8 \pm 5.9	29.3 \pm 4.6	25.1 \pm 4.9
PT _{CL}	0.13M	28.1 \pm 2.9	30.9 \pm 5.4	43.1 \pm 4.9
ModularPT	0.13M	43.2 \pm 0.6	44.8 \pm 4.9	61.8 \pm 1.8

Table 1: **Stage-agnostic** performance on News Classification, NER and Relation Extraction (RE). **Size** denotes average number of tunable parameters per training stage

explicitly imposes subset-invariant loss to adapt to dynamic label spaces, further boosting stage-agnostic performance by 14%-18% over PT_{CL}.

4.6 Results on Stage-fused Setting

We present results on our novel stage-fused setting in Table 2. We observe that none of the baselines are capable of handling this setting, as is evident from their abysmal performance across all testing stages. In absence of any label-modular representation, they are unable to utilize any information about the desired label-space. On the other hand, MODULARPROMPT not only outperforms all baselines by an average margin of 37.5%, it also achieves 4%-14% better performance than the oracle multi-task MODEL TUNING on News Classification and NER.

We select the top performing baselines in this setting and apply constrained decoding to them (marked with *), which improves their performance by 20%-30% on News and RE, 2%-4% on NER. However, MODULARPROMPT still outperforms these baselines by 14%-27%. This significant improvement is evident of the fact that MODULARPROMPT, by learning label-modular representations, can effectively combine partial knowledge from different training stages and condition the PLM on any target set of label prompts. This allows it to seamlessly adapt to dynamic unseen label spaces, without applying any post-processing or verbalizer.

Note that while PT_{CL} is able to combine knowledge from multiple training stages to support stage-agnostic testing, it fails to extract and consolidate specific knowledge corresponding to only the target label-set, across different stages.

Methods	Size	News	NER	RE
Finetune _{multitask}	770M	60.7 \pm 0.6	57.0 \pm 0.0	86.0 \pm 0.9
Finetune	770M	22.4 \pm 0.5	24.9 \pm 2.1	40.7 \pm 2.1
Finetune _{EWC}	770M	22.6 \pm 1.0	25.1 \pm 1.7	45.2 \pm 3.3
Finetune _{MAS}	770M	28.8 \pm 3.3	23.5 \pm 1.3	46.9 \pm 1.9
Adapter	4.8M	18.8 \pm 0.3	21.8 \pm 1.7	28.2 \pm 1.8
Adapter _{EWC}	4.8M	18.1 \pm 1.0	21.7 \pm 1.5	27.4 \pm 1.2
Adapter _{MAS}	4.8M	19.6 \pm 1.4	21.4 \pm 2.0	27.8 \pm 0.6
PT	0.13M	15.7 \pm 0.2	19.3 \pm 0.5	19.6 \pm 0.2
PT _{EWC}	0.13M	16.1 \pm 0.5	20.0 \pm 1.0	20.0 \pm 0.8
PT _{MAS}	0.13M	23.7 \pm 6.2	25.3 \pm 3.8	24.9 \pm 5.1
PT _{CL}	0.13M	27.5 \pm 2.8	30.4 \pm 5.1	43.3 \pm 5.2
Finetune _{MAS} *	770M	52.8 \pm 2.2	25.8 \pm 1.7	68.1 \pm 1.3
PT _{CL} *	0.13M	57.5 \pm 5.2	34.3 \pm 2.4	65.3 \pm 2.7
ModularPT	0.13M	74.8 \pm 1.7	61.6 \pm 2.2	82.4 \pm 0.8

Table 2: **Stage-fused** performance on News Classification, NER and relation extraction (RE)

4.7 Results on Stage-specific Setting

While MODULARPROMPT has proved to be particularly successful in handling the challenging non-stationary settings of stage-agnostic and stage-fused evaluations, we now want to see how competitive it is under stage-specific settings. From the results in Table 3, we see that the average stage-specific performance of MODULARPROMPT is comparable to vanilla PROMPT TUNING on the three tasks. Note that while MAS regularization boosts stage-agnostic performance somewhat for MODEL TUNING and PROMPT TUNING, it infact degrades their stage-specific performance by 10%-40%. Similarly applying EWC regularization fails to improve over the vanilla models in this setting while also proving less effective on stage-agnostic evaluation. This shows the lack of robustness of these techniques across the different non-stationary settings. But MODULARPROMPT is able to achieve state-of-the-art in stage-agnostic and stage-fused settings while remaining comparable to PROMPT TUNING in stage-specific evaluation. Besides, (Lester et al., 2021) showed that the performance gap between PROMPT TUNING and MODEL TUNING will gradually close as the size of backbone PLMs scales up. We posit that MODULARPROMPT, being an extension of PROMPT TUNING can similarly benefit from scaling-up of the PLM, but we leave this as future work owing to resource limitations.

Methods	Size	News	NER	RE
Finetune _{multitask}	770M	60.6 \pm 5.0	64.5 \pm 0.2	87.9 \pm 0.6
Finetune	770M	83.1 \pm 0.4	77.9 \pm 0.4	94.5 \pm 0.7
Finetune _{EWC}	770M	82.5 \pm 0.4	77.3 \pm 0.4	94.5 \pm 0.3
Finetune _{MAS}	770M	59.7 \pm 9.8	60.4 \pm 3.8	74.0 \pm 6.5
Adapter	4.8M	81.4 \pm 0.5	77.4 \pm 0.4	94.6 \pm 0.5
Adapter _{EWC}	4.8M	81.5 \pm 0.4	77.3 \pm 0.4	94.8 \pm 0.8
Adapter _{MAS}	4.8M	81.0 \pm 0.2	77.2 \pm 0.5	94.9 \pm 0.6
PT	0.13M	79.7 \pm 0.3	75.4 \pm 0.2	94.3 \pm 0.9
PT _{EWC}	0.13M	79.9 \pm 0.6	74.8 \pm 0.6	94.5 \pm 0.6
PT _{MAS}	0.13M	49.6 \pm 4.8	44.4 \pm 2.3	84.3 \pm 3.8
ModularPT	0.13M	80.4 \pm 0.6	74.5 \pm 0.6	93.6 \pm 0.4

Table 3: **Stage-specific** performance on News Classification, NER and RE (averaged over stages)

Methods	Stage-agnostic			Stage-fused		
	News	NER	RE	News	NER	RE
MODULARPROMPT	43.2 \pm 0.6	44.8 \pm 4.9	61.8 \pm 1.8	74.8 \pm 1.7	61.6 \pm 2.2	82.4 \pm 0.8
w/o transfer	35.4 \pm 2.7	43.0 \pm 3.6	54.8 \pm 2.7	64.9 \pm 1.7	57.4 \pm 2.3	77.5 \pm 1.2
w/o subset-inv	41.1 \pm 2.4	40.0 \pm 4.6	58.2 \pm 2.6	58.5 \pm 3.0	44.1 \pm 2.9	47.7 \pm 4.1

Table 4: **Ablation study** of MODULARPROMPT: average performance on stage-agnostic and stage-fused settings

Methods	News	NER	RE
ModularPT	74.8 \pm 1.7	61.6 \pm 2.2	82.4 \pm 0.8
drop ground-truth label prompt	1.5 \pm 0.5	4.1 \pm 0.8	0.5 \pm 0.4
drop one random label prompt	74.9 \pm 1.4	61.8 \pm 2.1	83.2 \pm 0.6
permute label prompt order	72.4 \pm 1.3	61.5 \pm 2.3	82.2 \pm 0.9

Table 5: Mean Stage-fused performance for different inference schemes

4.8 Ablation Study

We now analyze the contribution of different components of MODULARPROMPT towards its SoTA performance. From the results in Table 4, we see that in stage-agnostic setting, both label prompt transfer and subset invariant loss provide a boost, though the role of the former is seemingly more significant. On the contrary, removing subset invariant loss has a more debilitating effect on stage-fused performance. This evinces that subset invariant loss is indeed critical in learning label modular representations. This is essential to the stage-fused evaluation which needs to extract and dynamically re-compose label-specific knowledge.

4.9 Quantitative Analysis

Apart from achieving SoTA, does MODULARPROMPT possess the desirable characteristics of a modular model? According to Algorithm 1, MODULARPROMPT set $S = \Omega_i^{t,s}$ during inference. We ex-

Test stage 1: person {artist or author, ...}, organization {company, sports ...}
Input: person_author ... org_company ... Gilbert was a delegate to the <u>New York State Constitutional Convention</u> of 1894 .
Predict: Gilbert person politician ; New York State Constitutional Convention organization government ;
Input: person_author ... org_company ... After Bert Bell 's death in October , the 33-year-old <u>Rozelle</u> was the surprise choice for his replacement as <u>NFL</u> commissioner .
Predict: Bert Bell person athlete ; Rozelle person athlete ; NFL organization sports league ;

Figure 3: **Successful** (blue) and **failure** (red) cases of MODULARPROMPT predictions for stage-fused NER

periment with different strategies of input prompt construction including dropping label prompt(s) either corresponding to ground truth label(s) or one other random label, and permuting the default order of label prompts; see Table 5 for the results.

Indeed we observe that dropping the ground truth label prompt during inference degrades the mean performance by 57%-82% while dropping any other random label prompt boosts performance slightly. This strongly demonstrates the *label grounding* property of MODULARPROMPT, i.e. the knowledge of a class label is exclusively embedded in its corresponding label prompt. MODULARPROMPT also shows low sensitivity to the order of label prompts during inference - a yet another favourable property of label modular models

4.10 Qualitative Analysis

Revisiting Figure 1 presented in §1, we observe that MODULARPROMPT is able to predict correctly on a testing regime that is unseen during training, by extracting and consolidating label specific knowledge from multiple training stages. More example predictions are shown in Figure 3 (and Appendix B), which indicate that MODULARPROMPT is able to exploit in-context learning over label-prompts to generalize to unseen label-combinations during inference. For example, MODULARPROMPT tags “Gilbert” as politician as he was “a delegate to” a government. In the same spirit, MODULARPROMPT wrongly tags “Bert Bell” and “Rozelle” as athletes (true label being *person_other*) because they are associated with the sports league “NFL”. Such qualitative findings demonstrate MODULARPROMPT’s capabilities to learn label modular representations and integrate them dynamically during inference.

5 Conclusion

In this paper, we have proposed MODULARPROMPT, a novel label modular prompt tuning framework

for text classification *in-the-wild*. Extensive experiments show that MODULARPROMPT is able to consolidate knowledge learned during sequential training stages for stage-agnostic testing and extract and recompose knowledge for stage-fused testing, while maintaining competitive performance in stage-specific settings. We have also conducted analysis to show that MODULARPROMPT has desirable modular properties of label grounding, low order sensitivity and in-context learning. Being the first work on modular parameter efficient tuning, we hope for it to spur more research in this area in future towards solving a wider range of tasks under more general non-stationary settings.

Limitations

In this section, we discuss limitations and potential future work towards extending MODULARPROMPT to a more generalised method for wider applicability.

On Scalability In MODULARPROMPT, the input prompt \mathcal{T} grows in proportion to $|S|$, the size of label set of interest. This limits MODULARPROMPT from supporting huge label set (e.g., thousands of labels) as transformers can only condition on a bounded-length context. With long range transformers like Longformer (Beltagy et al., 2020), Performer (Choromanski et al., 2021) and LongT5 (Guo et al., 2021) coming into vogue, this issue is somewhat mitigated. Regardless of that, one potential solution is to formulate a hierarchical version of MODULARPROMPT, which is similar in spirit to hierarchical softmax (Morin and Bengio, 2005). Hierarchical MODULARPROMPT takes multiple steps for prediction, with each step to predict labels in a specific hierarchy level.

Another potential solution is to treat all label prompts as memory units from which the model learns to select relevant ones for a given data instance, in the spirit of (Wu et al., 2022).

On generation tasks AS MODULARPROMPT SHOWS SoTA performance and good modular characteristics for text classification *in-the-wild*, it is appealing to extend it to other tasks like Question Answering (QA), Machine Reading Comprehension (MRC) and Summarization. However, it is non-trivial for MODULARPROMPT to incorporate these tasks as their target texts are unstructured without clear class labels. One potential solution is to instead consider attributes or properties of target texts, which are also conditioning factors (e.g., formality, concise-

ness, topics, aspects, sentiment for summarization). With such definitions, it will be interesting to check if MODULARPROMPT framework can achieve good generalisation and conditional generation on text generation *in-the-wild*.

References

- Ferran Alet, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. 2018. Modular meta-learning. In *CoRL*, volume 87 of *Proceedings of Machine Learning Research*, pages 856–868. PMLR.
- 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.
- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *CVPR*, pages 39–48. IEEE Computer Society.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *CoRR*, abs/2004.05150.
- RC Berwick, AD Friederici, N Chomsky, and JJ Bolhuis. 2013. *Evolution, brain, and the nature of language*. *Trends Cogn Sci.*, 17(2).
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. Autoregressive entity retrieval. In *ICLR*. OpenReview.net.
- Ignacio Cases, Clemens Rosenbaum, Matthew Riemer, Atticus Geiger, Tim Klinger, Alex Tamkin, Olivia Li, Sandhini Agarwal, Joshua D. Greene, Dan Jurafsky, Christopher Potts, and Lauri Karttunen. 2019. Recursive routing networks: Learning to compose modules for language understanding. In *NAACL-HLT (1)*, pages 3631–3648. Association for Computational Linguistics.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. 2019. Efficient lifelong learning with A-GEM. In *ICLR (Poster)*. OpenReview.net.
- Tianqi Chen, Ian J. Goodfellow, and Jonathon Shlens. 2016. Net2net: Accelerating learning via knowledge transfer. In *ICLR*.
- Xinyun Chen, Chen Liang, Adams Wei Yu, Denny Zhou, Dawn Song, and Quoc V. Le. 2020. Neural symbolic reader: Scalable integration of distributed and symbolic representations for reading comprehension. In *ICLR*. OpenReview.net.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás Szepesvári, Peter Hawkins, Jared Quincy Davis, Afroz

- Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J. Colwell, and Adrian Weller. 2021. Rethinking attention with performers. In *ICLR*. OpenReview.net.
- Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein, and Trevor Darrell. 2021. Modular networks for compositional instruction following. In *NAACL-HLT*, pages 1033–1040. Association for Computational Linguistics.
- Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. 2019. Episodic memory in lifelong language learning. In *NeurIPS*, pages 13122–13131.
- Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Haitao Zheng, and Zhiyuan Liu. 2021. Few-nerd: A few-shot named entity recognition dataset. In *ACL/IJCNLP (1)*, pages 3198–3213. Association for Computational Linguistics.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *CoRR*, abs/2101.03961.
- Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. 2021. Recurrent independent mechanisms. In *ICLR*. OpenReview.net.
- Mandy Guo, Joshua Ainslie, David C. Uthus, Santiago Ontañón, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2021. Long5: Efficient text-to-text transformer for long sequences. *CoRR*, abs/2112.07916.
- Nitish Gupta, Kevin Lin, Dan Roth, Sameer Singh, and Matt Gardner. 2020. Neural module networks for reasoning over text. In *ICLR*. OpenReview.net.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. Fewrel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. In *EMNLP*, pages 4803–4809. Association for Computational Linguistics.
- 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 *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *CoRR*, abs/2106.09685.
- Ronghang Hu, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2018. Explainable neural computation via stack neural module networks. In *ECCV (7)*, volume 11211 of *Lecture Notes in Computer Science*, pages 55–71. Springer.
- Shengding Hu. 2022. Opendelta. <https://github.com/thunlp/OpenDelta>.
- Yichen Jiang and Mohit Bansal. 2019. Self-assembling modular networks for interpretable multi-hop reasoning. In *EMNLP/IJCNLP (1)*, pages 4473–4483. Association for Computational Linguistics.
- Tushar Khot, Daniel Khoshdel, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2021. Text modular networks: Learning to decompose tasks in the language of existing models. In *NAACL-HLT*, pages 1264–1279. Association for Computational Linguistics.
- 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.
- Louis Kirsch, Julius Kunze, and David Barber. 2018. Modular networks: Learning to decompose neural computation. In *NeurIPS*, pages 2414–2423.
- Sneha Kudugunta, Yanping Huang, Ankur Bapna, Maxim Krikun, Dmitry Lepikhin, Minh-Thang Luong, and Orhan Firat. 2021. Beyond distillation: Task-level mixture-of-experts for efficient inference. In *EMNLP (Findings)*, pages 3577–3599. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. In *EMNLP (1)*, pages 3045–3059. Association for Computational Linguistics.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *ACL/IJCNLP (1)*, pages 4582–4597. Association for Computational Linguistics.
- Daqing Liu, Hanwang Zhang, Zheng-Jun Zha, and Feng Wu. 2019. Learning to assemble neural module tree networks for visual grounding. In *ICCV*, pages 4672–4681. IEEE.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. *GPT understands, too*. *CoRR*, abs/2103.10385.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *NIPS*, pages 6467–6476.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021a. Compacter: Efficient low-rank hypercomplex adapter layers. In *NeurIPS*, pages 1022–1035.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021b. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. In *ACL/IJCNLP (1)*, pages 565–576. Association for Computational Linguistics.

- Arun Mallya, Dillon Davis, and Svetlana Lazebnik. 2018. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *ECCV (4)*, volume 11208 of *Lecture Notes in Computer Science*, pages 72–88. Springer.
- Rishabh Misra. 2018. [News category dataset](#).
- Frederic Morin and Yoshua Bengio. 2005. Hierarchical probabilistic neural network language model. In *AIS-TATS*. Society for Artificial Intelligence and Statistics.
- Oleksiy Ostapenko, Pau Rodríguez, Massimo Caccia, and Laurent Charlin. 2021. Continual learning via local module composition. In *NeurIPS*, pages 30298–30312.
- Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *EACL*, pages 487–503. Association for Computational Linguistics.
- Edoardo Maria Ponti, Ivan Vulic, Ryan Cotterell, Marinela Parovic, Roi Reichart, and Anna Korhonen. 2021. Parameter space factorization for zero-shot learning across tasks and languages. *Trans. Assoc. Comput. Linguistics*, 9:410–428.
- 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.
- Guanghai Qin and Jason Eisner. 2021. Learning how to ask: Querying lms with mixtures of soft prompts. In *NAACL-HLT*, pages 5203–5212. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Janarthanan Rajendran, Aravind S. Lakshminarayanan, Mitesh M. Khapra, P. Prasanna, and Balaraman Ravindran. 2017. Attend, adapt and transfer: Attentive deep architecture for adaptive transfer from multiple sources in the same domain. In *ICLR (Poster)*. OpenReview.net.
- Clemens Rosenbaum, Ignacio Cases, Matthew Riemer, and Tim Klinger. 2019. Routing networks and the challenges of modular and compositional computation. *CoRR*, abs/1904.12774.
- Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. 2019. Latent multi-task architecture learning. In *AAAI*, pages 4822–4829. AAAI Press.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. 2016. Progressive neural networks. *CoRR*, abs/1606.04671.
- Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. 2018. Progress & compress: A scalable framework for continual learning. In *International Conference on Machine Learning*, pages 4528–4537. PMLR.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc V. Le, Geoffrey E. Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR (Poster)*. OpenReview.net.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Fan-Keng Sun, Cheng-Hao Ho, and Hung-Yi Lee. 2020. LAMOL: language modeling for lifelong language learning. In *ICLR*. OpenReview.net.
- Sebastian Thrun. 1996. [Is learning the n-th thing any easier than learning the first?](#) In *Advances in Neural Information Processing Systems*, volume 8. MIT Press.
- Tu Vu, Brian Lester, Noah Constant, Rami Al-Rfou’, and Daniel Cer. 2022. Spot: Better frozen model adaptation through soft prompt transfer. In *ACL (1)*, pages 5039–5059. Association for Computational Linguistics.
- Yuhuai Wu, Markus N. Rabe, DeLesley Hutchins, and Christian Szegedy. 2022. Memorizing transformers. *CoRR*, abs/2203.08913.
- Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. 2022. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *ACL (2)*, pages 1–9. Association for Computational Linguistics.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 3987–3995. PMLR.
- Qi Zhu, Bing Li, Fei Mi, Xiaoyan Zhu, and Minlie Huang. 2022. Continual prompt tuning for dialog state tracking. In *ACL (1)*, pages 1124–1137. Association for Computational Linguistics.

A More Details on Dataset Construction

In this section, we provide more complementary details to §4.1. Table 6 shows examples of text input and target for three datasets. For Huffpost News and FewRel, we subsample exactly 200,50 examples per label class. For FewNERD, as it can have multiple label types per example, we subsample at 50 examples per label class.

Dataset	Example
Huffpost	'Tis The Season To Be Cheeky With 'Jingle Butts' Music Video. Clip features both male and female butt models. → weird news
	The 4 Best Last-Minute Christmas Gift Ideas. Scrambling to shop for that last relative, significant other or hard-to-buy-for friend? We're here to help. → style and beauty
FewNERD	One of the Yak-3s was destroyed right away . → Yak-3s ! product airplane ;
	All Together Now is a British reality television music competition which first aired on BBC One on 27 January 2018 . → All Together Now ! event other ; BBC One ! organization media ;
FewRel	Its main base was at Tampere - Pirkkala Airport (TMP) , Tampere . Pirkkala Airport Tampere → place served by transport hub
	He represented Sweden at the 1934 FIFA World Cup . 1934 FIFA World Cup Sweden → participating teams

Table 6: Examples of text input → text target for Huffpost News, fewNERD and FewRel

B More Qualitative Examples

Similar to §4.10, we show another example of MODULARPROMPT on stage-fused NER in Figure 4 and more example predictions in Figure 5. These additional examples strengthen the conclusions in §4.10. In the second example in Figure 5, MODULARPROMPT tags "Big Twin Sauce" as a *product food* while its ground truth tag is *product other*. We can see MODULARPROMPT considers the context as the entity is associated with a restaurant. Similarly, in the third example, "Kobo Touch" is actually a hardware reader and its ground truth tag is *product other*. However, such world knowledge is not available and MODULARPROMPT tags it as a software based on the context of "eBooks" and libraries.

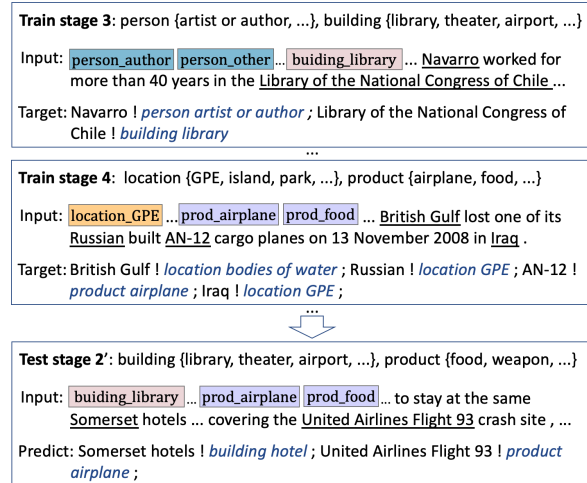


Figure 4: Another example of MODULARPROMPT for stage-fused NER



Figure 5: Successful (blue) and failure (red) cases of MODULARPROMPT predictions for stage-fused NER

C Detailed Results on Stage-agnostic and Stage-specific Settings

Table 7 shows detailed stage-agnostic results. $D_{\{1..k\}}$ corresponds to multi-stage training over domains $\{1, \dots, k\}$ and testing over the aggregated domains. For baselines, we use their checkpoints after k-th stage for evaluating $D_{\{1..k\}}$. For MODULARPROMPT, we use the final checkpoint. Table 8, 9 show detailed stage-fused results. D_k denotes k-th task fused stage.

Methods	News Classification				NER			RE			
	D _{1..2}	D _{1..3}	D _{1..4}	D _{1..5}	D _{1..2}	D _{1..3}	D _{1..4}	D _{1..2}	D _{1..3}	D _{1..4}	D _{1..5}
Finetune _{Multitask}	60.0±5.2	61.6±4.8	62.1±4.3	60.6±5.0	64.3±0.3	68.1±0.3	64.5±0.2	85.3±1.0	87.4±0.6	83.7±2.2	87.9±0.6
Finetune	57.8±0.9	35.2±1.8	29.4±1.2	23.0±0.6	39.8±0.4	30.9±2.1	25.0±2.7	71.1±2.2	57.4±2.5	48.0±1.6	40.8±2.2
Finetune _{EWC}	61.4 ±2.6	35.3±1.8	27.3±2.5	22.5±1.2	42.9±2.2	37.8±4.8	26.6±2.8	79.2±3.5	61.2±6.5	52.5±3.4	45.4±3.2
Finetune _{MAS}	58.9±1.1	46.9±2.8	38.7±3.0	29.2±2.9	48.6±7.5	37.1±2.1	25.4±2.9	74.5±3.9	61.3±2.7	54.7±1.5	46.2±1.8
Adapter	45.2±2.6	29.3±0.4	24.5±0.5	19.3±0.3	38.9±1.0	27.7±0.7	21.2±2.0	56.6±1.4	40.4±2.5	36.9±1.2	28.6±1.9
Adapter _{EWC}	45.6±3.0	29.6±0.5	24.4±0.6	18.6±1.0	38.6±0.7	29.1±1.7	20.5±1.3	54.9±1.2	42.6±1.5	36.1±0.8	27.8±1.2
Adapter _{MAS}	45.6±2.6	30.0±0.8	26.1±2.5	20.1±1.5	38.2±0.3	28.6±1.1	21.1±2.0	57.1±2.9	40.5±1.9	35.5±1.9	28.3±0.5
PT	36.9±0.5	27.9±0.4	21.7±0.4	16.1±0.3	37.0±0.3	26.1±0.3	18.8±0.1	47.3±1.0	32.4±0.5	24.3±0.4	19.6±0.1
PT _{EWC}	36.9±1.1	28.2±0.2	21.5±0.3	16.5±0.5	36.7±0.8	26.0±0.4	19.0±0.5	46.8±0.5	32.3±1.2	25.4±0.9	20.1±0.8
PT _{MAS}	43.0±3.5	30.1±5.2	24.1±4.5	23.8±5.9	28.8±2.8	29.2±6.0	29.3±4.6	51.3±1.3	36.0±0.9	30.7±1.7	25.1±4.9
PT _{CL}	60.3±3.7	41.8±7.5	35.4±3.3	28.1±2.9	47.4±4.6	39.7±3.9	30.9±5.4	69.9±7.8	54.3±3.9	47.9±4.2	43.1±4.9
ModularPT	60.3±3.9	53.9 ±2.3	51.1 ±1.4	43.2 ±0.6	63.3 ±1.4	53.9 ±4.3	44.8 ±4.9	81.3 ±3.3	70.1 ±3.7	65.1 ±2.2	61.8 ±1.8
w/o transfer	61.2±2.8	48.9±2.4	43.1±2.7	35.4±2.7	62.6±1.4	53.8±1.7	43.0±3.6	82.2±2.3	67.2±3.2	59.4±2.0	54.8±2.7
w/o subset_inv	62.4±1.9	53.9±3.2	50.0±2.9	41.1±2.4	60.0±2.8	51.2±3.5	40.0±4.6	74.4±10.9	66.2±4.8	60.3±3.5	58.2±2.6

Table 7: Detailed Stage-agnostic Performance on News Classification, NER and RE.

Methods	News Classification						NER				
	D1'	D2'	D3'	D4'	D5'	Mean	D1'	D2'	D3'	D4'	Mean
Finetune _{Multitask}	58.5±0.7	57.5±1.7	60.8±2.7	64.4±1.8	62.2±1.6	60.7±0.6	64.3±1.0	53.8±0.1	51.0±0.7	59.1±0.2	57.0±0.0
Finetune	20.5±1.5	21.7±1.2	21.8±1.3	24.5±1.3	23.4±1.0	22.4±0.5	20.9±2.9	24.1±2.9	35.4±1.8	19.1±3.4	24.9±2.1
Finetune _{EWC}	20.9±1.7	21.9±0.7	21.2±2.2	25.6±1.8	23.3±1.7	22.6±1.0	21.8±1.5	24.6±1.4	34.7±2.0	19.3±4.2	25.1±1.7
Finetune _{MAS}	26.4±6.2	25.6±2.9	26.8±3.4	38.4±7.0	26.5±6.4	28.8±3.3	22.4±3.5	25.2±0.9	19.1±4.4	27.2±5.9	23.5±1.3
Adapter	15.8±0.9	20.5±0.4	15.5±2.0	22.6±0.6	19.6±1.2	18.8±0.3	17.1±2.0	21.6±3.3	33.5±1.7	15.0±1.9	21.8±1.7
Adapter _{EWC}	15.2±1.3	20.4±0.5	13.7±2.0	23.5±1.9	18.0±1.4	18.1±1.0	17.3±1.8	21.0±2.5	33.8±1.6	14.5±1.5	21.7±1.5
Adapter _{MAS}	17.3±2.0	20.2±0.4	16.5±1.3	23.5±2.3	20.3±2.5	19.6±1.4	16.7±2.1	20.4±2.5	34.1±2.3	14.2±2.0	21.4±2.0
PT	13.7±0.3	19.6±0.3	8.3±0.2	21.0±1.3	15.6±0.6	15.7±0.2	14.9±1.2	17.2±0.8	31.9±1.3	13.2±1.0	19.3±0.5
PT _{EWC}	13.7±0.4	19.7±0.2	9.2±1.5	21.9±0.5	15.8±1.1	16.1±0.5	14.7±0.8	18.5±1.9	33.0±1.3	14.0±1.5	20.0±1.0
PT _{MAS}	25.3±6.0	21.9±3.9	20.0±13.4	27.6±7.9	23.8±5.3	23.7±6.2	26.1±3.1	20.7±2.5	22.3±4.2	31.9±7.8	25.3±3.8
PT _{CL}	25.7±5.9	27.1±5.5	21.9±5.5	29.1±7.3	33.9±4.8	27.5±2.8	28.1±4.3	34.7±4.8	30.5±6.0	28.2±7.9	30.4±5.1
Finetune _{MAS} *	47.9±2.4	49.7±5.5	55.8±4.7	52.4±2.5	58.2±4.9	52.8±2.2	23.3±2.8	29.3±1.8	22.3±1.7	28.3±2.0	25.8±1.7
PT _{CL} *	56.1±6.3	62.5±4.7	57.5±6.6	55.8±9.4	55.5±2.0	57.5±5.2	33.1±4.0	37.4±1.4	28.2±4.9	38.6±2.0	34.3±2.4
ModularPT	71.6 ±3.0	75.8 ±4.0	76.5 ±2.0	79.1 ±3.3	71.0 ±1.6	74.8 ±1.7	62.8 ±4.0	60.7 ±1.1	61.5 ±2.5	61.5 ±4.4	61.6 ±2.2
w/o transfer	67.8±3.9	69.7±3.0	66.7±5.3	63.8±3.3	56.6±5.3	64.9±1.7	56.0±4.5	51.2±7.6	61.3±1.6	61.2±2.7	57.4±2.3
w/o subset_inv	61.3±7.7	55.9±11.8	60.4±5.2	61.0±3.0	53.7±7.2	58.5±3.0	51.9±9.5	37.5±3.7	40.7±7.3	46.3±4.3	44.1±2.9

Table 8: Detailed Stage-fused Performance on News Classification and NER

Methods	RE					Mean
	D1'	D2'	D3'	D4'	D5'	
Finetune _{Multitask}	80.5 \pm 2.1	90.3 \pm 1.6	85.4 \pm 1.2	88.5 \pm 0.9	85.4 \pm 1.0	86.0 \pm 0.9
Finetune	31.2 \pm 2.8	49.6 \pm 3.5	37.4 \pm 4.0	43.3 \pm 3.7	42.2 \pm 1.3	40.7 \pm 2.1
Finetune _{EWC}	34.3 \pm 4.1	57.8 \pm 3.6	40.6 \pm 2.5	46.9 \pm 6.0	46.5 \pm 2.1	45.2 \pm 3.3
Finetune _{MAS}	41.7 \pm 3.0	60.1 \pm 3.4	52.2 \pm 4.2	46.1 \pm 3.8	34.4 \pm 3.7	46.9 \pm 1.9
Adapter	21.0 \pm 0.6	33.8 \pm 4.6	21.3 \pm 1.3	32.9 \pm 2.0	32.0 \pm 2.6	28.2 \pm 1.8
Adapter _{EWC}	20.1 \pm 0.7	31.1 \pm 3.7	19.7 \pm 0.3	34.0 \pm 2.8	32.0 \pm 3.4	27.4 \pm 1.2
Adapter _{MAS}	20.1 \pm 0.5	32.2 \pm 2.1	22.5 \pm 1.6	31.6 \pm 3.1	32.7 \pm 2.4	27.8 \pm 0.6
PT	19.6 \pm 0.1	20.8 \pm 0.7	20.3 \pm 1.2	18.5 \pm 0.2	19.0 \pm 0.3	19.6 \pm 0.2
PT _{EWC}	19.8 \pm 0.1	22.2 \pm 3.0	19.9 \pm 0.4	18.5 \pm 0.1	19.8 \pm 1.1	20.0 \pm 0.8
PT _{MAS}	18.8 \pm 6.9	25.0 \pm 9.1	22.8 \pm 5.9	32.7 \pm 8.1	25.5 \pm 3.7	24.9 \pm 5.1
PT _{CL}	37.6 \pm 9.1	51.1 \pm 4.4	37.8 \pm 4.4	41.4 \pm 4.7	48.4 \pm 6.1	43.3 \pm 5.2
Finetune _{MAS} *	55.7 \pm 3.8	72.9 \pm 2.6	69.8 \pm 3.5	80.2 \pm 2.2	62.0 \pm 2.8	68.1 \pm 1.3
PT _{CL} *	58.8 \pm 5.1	65.1 \pm 4.0	61.1 \pm 1.6	78.3 \pm 3.3	63.3 \pm 3.5	65.3 \pm 2.7
ModularPT	80.3 \pm 1.6	89.2 \pm 2.5	75.1 \pm 3.9	88.5 \pm 2.7	78.6 \pm 0.8	82.4 \pm 0.8
w/o transfer	73.5 \pm 2.6	85.1 \pm 1.9	66.9 \pm 7.4	88.2 \pm 2.6	73.7 \pm 3.6	77.5 \pm 1.2
w/o subset_inv	45.1 \pm 5.1	50.6 \pm 6.1	48.3 \pm 5.9	46.5 \pm 4.8	48.0 \pm 6.7	47.7 \pm 4.1

Table 9: Detailed Stage-fused Performance on RE