

Self-Supervised Meta-Learning for Few-Shot Natural Language Classification Tasks

Trapit Bansal^{◇*} and Rishikesh Jha[†] and Tsendsuren Munkhdalai[‡] and Andrew McCallum[◇]

[◇]University of Massachusetts, Amherst

[†]Code for Science and Society

[‡]Microsoft Research, Montréal, Canada

Abstract

Self-supervised pre-training of transformer models has revolutionized NLP applications. Such pre-training with language modeling objectives provides a useful initial point for parameters that generalize well to new tasks with fine-tuning. However, fine-tuning is still data inefficient — when there are few labeled examples, accuracy can be low. Data efficiency can be improved by optimizing pre-training directly for future fine-tuning with few examples; this can be treated as a meta-learning problem. However, standard meta-learning techniques require many training tasks in order to generalize; unfortunately, finding a diverse set of such supervised tasks is usually difficult. This paper proposes a self-supervised approach to generate a large, rich, meta-learning task distribution from unlabeled text. This is achieved using a cloze-style objective, but creating separate multi-class classification tasks by gathering tokens-to-be blanked from among only a handful of vocabulary terms. This yields as many unique meta-training tasks as the number of subsets of vocabulary terms. We meta-train a transformer model on this distribution of tasks using a recent meta-learning framework. On 17 NLP tasks, we show that this meta-training leads to better few-shot generalization than language-model pre-training followed by finetuning. Furthermore, we show how the self-supervised tasks can be combined with supervised tasks for meta-learning, providing substantial accuracy gains over previous supervised meta-learning.

1 Introduction

Self-supervised learning has emerged as an important training paradigm for learning model parameters which are more generalizable and yield better representations for many down-stream tasks. This typically involves learning through labels that come

naturally with data, for example words in natural language. Self-supervised tasks typically pose a supervised learning problem that can benefit from lots of naturally available data and enable pre-training of model parameters that act as a useful prior for supervised fine-tuning (Erhan et al., 2010). Masked language modeling (Devlin et al., 2018), and other related approaches (Peters et al., 2018; Howard and Ruder, 2018; Radford et al., 2019), is an example of such a self-supervised task that is behind the success of transformer models like BERT.

While self-supervised pre-training is beneficial, it has been recently noted that it is not data-efficient and typically requires large amounts of fine-tuning data for good performance on a target task (Yogatama et al., 2019; Bansal et al., 2019). This can be evaluated as a few-shot learning problem, where a model is given only few examples of a new task and is expected to perform well on that task. This paper focuses on this problem of few-shot learning and develops models which demonstrate better few-shot generalization to new tasks.

Large scale pre-training suffers from a train-test mismatch as the model is not optimized to learn an initial point that yields good performance when fine-tuned with few examples. Moreover, fine-tuning of a pre-trained model typically introduces new random parameters, such as softmax layers, and important hyper-parameters such as learning rate, which are hard to estimate robustly from the few examples. Thus, we propose to remove this train-test mismatch, and treat learning an initial point and hyper-parameters jointly from unlabelled data, which allows data-efficient fine-tuning, as a meta-learning problem.

Meta-learning, or learning to learn (Thrun and Pratt, 2012; Schmidhuber, 1987), treats learning a parameterized algorithm, such as a neural net optimized with SGD, that generalizes to new tasks as a learning problem. This typically assumes access

*Correspondence: tbansal@cs.umass.edu

to a distribution over tasks in order to enable learning. Creating tasks which enable meta-learning is one of the main challenges for meta-learning (Bengio et al., 1992; Santoro et al., 2016; Vinyals et al., 2016), and typical supervised meta-learning approaches create task distributions from a fixed task dataset with large number of labels by sub-sampling from the set of labels (Vinyals et al., 2016; Ravi and Larochelle, 2017). While this enables generalization to new labels, it limits generalization to unseen tasks due to over-fitting to the training task distribution (Yin et al., 2020). Moreover, large supervised datasets with a large label set are not always available for meta-learning, as is often the case in many NLP applications.

To overcome these challenges of supervised meta-learning, we propose a self-supervised approach and create the task-distribution from unlabelled sentences. Taking inspiration from the cloze task (Taylor, 1953), we create separate multi-class classification tasks by gathering tokens-to-be blanked from a subset of vocabulary words, allowing for as many unique meta-training tasks as the number of subsets of words in the language. The proposed approach, which we call Subset Masked Language Modeling Tasks (SMLMT), enables training of meta-learning methods for NLP at a much larger scale than was previously feasible while also ameliorating the risk of over-fitting to the training task distribution. This opens up new possibilities for applications of meta-learning in NLP, such as few-shot learning, continual learning, architecture search and more.

This work focuses on few-shot learning and makes the following contributions: (1) we introduce a self-supervised approach to create tasks for meta-learning in NLP, Subset Masked Language Modeling Tasks (SMLMT), which enables application of meta-learning algorithms for goals like few-shot learning; (2) utilizing SMLMT as the training task distribution, we train a state-of-the-art transformer architecture, BERT (Devlin et al., 2018), using a recent optimization-based meta-learning method which was developed for diverse classification tasks (Bansal et al., 2019); (3) we show that the self-supervised SMLMT can also be combined with supervised task data to enable better feature learning, while still allowing for better generalization by avoiding meta-overfitting to the supervised tasks through the use of SMLMT; (4) we rigorously evaluate the proposed approach on few-shot gener-

alization to unseen tasks as well as new domains of tasks seen during training and show that the proposed approach demonstrates better generalization than self-supervised pre-training or self-supervised pre-training followed by multi-task training; (5) we also study the effect of number of parameters for few-shot learning and find that while bigger pre-trained or meta-trained models generalize better than smaller models, meta-learning leads to substantial gains even for the smaller models.

2 Preliminaries

In supervised meta-learning, we typically assume access to a task distribution $\mathcal{P}(\mathcal{T})$. Practically, this translates to a fixed set of training tasks $\{T_1, \dots, T_M\}$, which are referred to as meta-training tasks. For supervised classification, each task T_i is an N_i -way classification task. While many meta-learning algorithms assume a fixed N -way classification, we follow the more practical approach of Bansal et al. (2019) and allow for a diverse set of classification tasks with potentially different number of classes.

The goal of a meta-learning algorithm is to utilize the meta-training tasks to learn a learning procedure that generalizes to held-out tasks $T' \sim \mathcal{P}(\mathcal{T})$. Model-agnostic meta-learning (MAML) (Finn et al., 2017) is an example of such a meta-learning algorithm. MAML learns an initial point θ for a classifier $f_\theta : x \rightarrow \hat{y}$, that can be optimized via gradient descent on the supervised loss \mathcal{L}_i defined for the task T_i , using its support set $\mathcal{D}^{tr} \sim T_i$:

$$\theta'_i \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}_i(\mathcal{D}^{tr}, \theta) \quad (1)$$

where α is the learning rate. The optimized point θ' is then evaluated on another validation set for the task, $\mathcal{D}^{val} \sim T_i$, using the loss function \mathcal{L}_i . This loss across meta-training tasks serves as the training error to optimize the initial point and parameters like learning-rate ($\Theta := \{\theta, \alpha\}$):

$$\Theta \leftarrow \Theta - \beta \nabla_{\Theta} \mathbb{E}_{T_i \sim \mathcal{P}(\mathcal{T})} \left[\mathcal{L}_i(\mathcal{D}^{val}, \theta'_i) \right] \quad (2)$$

where β is the learning rate for the meta-training process. Training proceeds in an episodic framework (Vinyals et al., 2016), where in each episode a mini-batch of tasks are sampled along with their support and validation sets, and the model parameters are optimized using (1) and (2), which are also referred to as inner and outer loop, respectively.

Meta-training Tasks: We summarize how supervised task data-sets are typically leveraged to

Subset: {Democratic, Capital}

↓

Support set	
Sentence	Class
A member of the [m] Party, he was the first African American to be elected to the presidency.	1
The [m] Party is one of the two major contemporary political parties in the United States, along with its rival, the Republican Party.	1
Honolulu is the [m] and largest city of the U.S. state of Hawaii.	2
Washington, D.C., formally the District of Columbia and commonly referred to as Washington or D.C., is the [m] of the United States.	2

Query: New Delhi is an urban district of Delhi which serves as the [m] of India
Correct Prediction: 2

Figure 1: An example of a 2-way 2-shot task in SMLMT. The support set and one query is shown. Any N -way k -shot task can be constructed similarly.

create meta-training tasks (Vinyals et al., 2016). Assuming access to a supervised task with L classes, an N -way k -shot task is created by first sampling N classes, assuming $N \ll L$. Then for each of the N sampled classes, $(k + q)$ examples of each class is randomly sampled from the dataset and assigned a unique label in $\{1, \dots, N\}$. The k examples for each label serve as the support set, while the q examples constitute the validation set described above. Note, that each task consists of a small subset of classes and the class to label (1 to N) assignment is random. This is crucial to avoid learning the sample to label bindings in the parameters of the model, which will make the task-specific training (in (1)) irrelevant and the model will not generalize to new tasks. An example of this approach is MiniImageNet (Ravi and Larochelle, 2017), which is a benchmark dataset for learning few-shot image classification.

3 Self-supervised Tasks for Meta-learning

The existing approach to using a supervised dataset to create tasks, as described above, is fraught with issues, specially for NLP applications. First, note that large classification datasets with large label spaces are not readily available for all NLP tasks, for example sentiment classification which has only few discrete labels. Second, limiting to a fixed supervised dataset to create tasks limits generalization ability and the meta-learned models might generalize to new labels for the task but fail to generalize to new novel tasks (Metz et al., 2019). Lastly, such an approach is also not feasible in all problems where we will like to apply meta-learning (Yin et al., 2020). For example, consider meta-learning a natural language inference (NLI) model across multiple domains which can generalize to

new domains. A powerful model can ignore the training data for each task and directly learn to predict the NLI tag for the examples in each training domain, which will lead to low training error but the model will not generalize to new domains. We overcome these issues by utilizing unlabelled data to create meta-learning tasks. See Fig. 1 for an example of generated task.

Subset Masked Language Modeling Tasks (SMLMT):

We are given a text corpus split into sentences X_i and each sentence is a sequence of words from a vocabulary of size V . Now, in Subset Masked Language Modeling Tasks, each task is defined from a *subset* of vocabulary words. To create an N -way classification task, we randomly select N unique vocabulary words: $\{v_1, \dots, v_N\}$. Then we consider all sentences containing these N words, and for each word randomly sample $r = k + q$ sentences: $\mathbf{x}_{v_i} = \{X_1, \dots, X_r | v_i \in X_i\}$. Now, we mask the corresponding chosen word from the sentences in each of these N sets, so $\mathbf{x}'_{v_i} = \{\text{Mask}(X_1, v_i), \dots, \text{Mask}(X_r, v_i)\}$ where $\text{Mask}(X, v)$ replaces all occurrences of v in X with the mask token $[m]$. The set $\{\mathbf{x}'_{v_1}, \dots, \mathbf{x}'_{v_N}\}$ is then a well-defined N -partition of $N \times r$ sentences, that serves as input examples for the N -way classification task. We forget the original word corresponding to the masked tokens in these sets and assign labels in $\{1, \dots, N\}$ to the N sets. This gives an instance of an SMLMT classification task: $T = \{(x_{ij}, i) | i \in \{1, \dots, N\}, x_{ij} \in \mathbf{x}'_{v_i}\}$. This can be split into support and validation for meta-training.

In an SMLMT instance, each input sentence consists of exactly one word that is masked throughout the sentence and its label corresponds to that word. This requires a similar reasoning ability as cloze tasks (Taylor, 1953). Moreover, crucially, the SMLMT task creation ensures that a model cannot memorize the input-label mapping as the target masked word is hidden and the label assignment is randomized, requiring the model to infer the labels from the support set. Note that the SMLMT tasks are also closely related to masked language modeling (MLM) (Devlin et al., 2018). While MLM is a word-level classification task, SMLMT is a sentence-level classification task. Each unique subset of words from the vocabulary defines a unique task in SMLMT. This allows for as many unique tasks as the *number of subsets of words in the vocabulary*, enabling large-scale meta-learning from unsupervised data.

Hybrid SMLMT: Tasks from SMLMT can also be combined with supervised tasks to encourage better feature learning (Caruana, 1997) and increase diversity in tasks for meta-learning. We use a sampling ratio $\lambda \in (0, 1)$ and in each episode select an SMLMT task with probability λ or a supervised task with probability $(1 - \lambda)$. The use of SMLMT jointly with supervised tasks ameliorates meta-overfitting, as tasks in SMLMT cannot be solved without using the task support data. λ is a hyper-parameter. In our experiments, we found $\lambda = 0.5$ to work well.

4 Meta-learning Model

We now discuss the meta-learning model for learning new NLP tasks.

Text encoder: The input to the model is natural language sentences. This is encoded using a transformer (Vaswani et al., 2017) text encoder. We follow the BERT (Devlin et al., 2018) model and use the same underlying neural architecture for the transformer as their base model. Given an input sentence, the transformer model yields contextualized token representations for each token in the input after multiple layers of self-attention. Following BERT, we add a special CLS token to the start of the input that is used as a sentence representation for classification tasks. Given an input sentence X , let $f_\pi(X)$ be the CLS representation of the final layer of the transformer with parameters π .

Meta-learning across diverse classes: Our motivation is to meta-learn an initial point that can generalize to novel NLP tasks, thus we consider methods that apply to diverse number of classes. Note that many meta-learning models only apply to a fixed number of classes (Finn et al., 2017) and require training different models for different number of classes. We follow the approach of Bansal et al. (2019) that learns to generate softmax classification parameters conditioned on a task support set to enable training meta-learning models that can adapt to tasks with diverse classes. This combines benefits of metric-based methods (Vinyals et al., 2016; Snell et al., 2017) and optimization-based methods for meta-learning. The key idea is to train a deep set encoder $g_\psi(\cdot)$, with parameters ψ , which takes as input the set of examples of a class n and generates a $(d+1)$ dimensional embedding that serves as the linear weight and bias for class n in the softmax classification layer. Let $\{X_{1n}, \dots, X_{kn}\}$ be the k

examples for class n in the support set of a task t :

$$w_t^n, b_t^n = g_\psi(\{f_\pi(X_{1n}), \dots, f_\pi(X_{kn})\}) \quad (3)$$

$$p(y|X) = \text{softmax}\{\mathbf{W}_t h_\phi(f_\pi(X)) + \mathbf{b}_t\} \quad (4)$$

where $\mathbf{W}_t = [w_t^1; \dots; w_t^N] \in \mathcal{R}^{N \times d}$, $\mathbf{b}_t = [b_t^1; \dots; b_t^N] \in \mathcal{R}^d$ are the concatenation of the per-class vectors in (3), and h_ϕ is a MLP with parameters ϕ and output dimension d .

Using the above model to generate predictions, the parameters are meta-trained using the MAML algorithm (Finn et al., 2017). Concretely, set $\theta := \{\pi, \phi, \mathbf{W}_t, \mathbf{b}_t\}$ for the task-specific inner loop gradient updates in (1) and set $\Theta := \{\pi, \psi, \alpha\}$ for the outer-loop updates in (2). Note that we do multiple steps of gradient descent in the inner loop. Bansal et al. (2019) performed extensive ablations over parameter-efficient versions of the model and found that adapting all parameters with learned per-layer learning rates performs best for new tasks. We follow this approach. Full training algorithm can be found in the Appendix.

Fast adaptation: Flennerhag et al. (2019) proposed an approach which mitigates slow adaption often observed in MAML by learning to warp the task loss surface to enable rapid descent to the loss minima. This is done by interleaving a neural network’s layers with non-linear layers, called warp layers, which are not adapted for each task but are still optimized across tasks in the outer-loop updates in (2). Since introducing additional layers will make computation more expensive, we use existing transformer layers as warp layers. We designate the feed-forward layers in between self-attention layers of BERT, which project from dimension 768 to 3072 to 768, as warp-layers. Note that these parameters also constitute a large fraction of total parameters ($\sim 51\%$). Thus in addition to the benefit from warping, not adapting these layers per task means significantly faster training and smaller number of per-task parameters during fine-tuning. The warp layers are still updated in the outer loop during meta-training.

5 Related Work

Language model pre-training has recently emerged as a prominent approach to learning general purpose representations (Howard and Ruder, 2018; Peters et al., 2018; Devlin et al., 2018; Radford et al., 2019; Yang et al., 2019; Raffel et al., 2019). Refer to Weng (2019) for a review of self-supervised learning. Pre-training is usually a two step process and fine-tuning introduces random parameters

making it inefficient when target tasks have few examples (Bansal et al., 2019). Multi-task learning of pre-trained models has shown improved results on many tasks (Phang et al., 2018; Liu et al., 2019a). More recently, and parallel to this work, Brown et al. (2020) show that extremely large language models can act as few-shot learners. They propose a query-based approach where few-shot task data is used as context for the language model. In contrast, we employ a fine-tuning based meta-learning approach that enjoys nice properties like consistency which are important for good out-of-distribution generalization (Finn, 2018). Moreover, we show that self-supervised meta-learning can also improve few-shot performance for smaller models.

Meta-Learning methods can be categorized as: optimization-based (Finn et al., 2017; Li et al., 2017; Nichol and Schulman, 2018; Rusu et al., 2018), model-based (Santoro et al., 2016; Ravi and Larochelle, 2017; Munkhdalai and Yu, 2017), and metric-based (Vinyals et al., 2016; Snell et al., 2017). Refer to Finn (2018) for an exhaustive review. Unsupervised meta-learning has been explored in vision. Hsu et al. (2019) cluster images using pre-trained embeddings to create tasks for meta-learning. Metz et al. (2019) meta-learn a biologically-motivated update rule from unsupervised data in a semi-supervised framework. Compared to these, we directly utilize text data to automatically create unsupervised tasks without relying on pre-trained embeddings or access to target tasks.

In NLP, meta-learning approaches have followed the recipe of using supervised task data and learning models for specific tasks. Such approaches (Yu et al., 2018; Gu et al., 2018; Guo et al., 2018; Han et al., 2018; Mi et al., 2019) train to generalize to new labels of a specific task like relation classification and don't generalize to novel tasks. Bansal et al. (2019) proposed an approach that applies to diverse tasks to enable practical meta-learning models and evaluate on generalization to new tasks. However, they rely on supervised task data from multiple tasks and suffer from meta-overfitting as we show in our empirical results. Holla et al. (2020) studied related approaches for the task of word-sense disambiguation. To the best of our knowledge, the method proposed here is the first self-supervised approach to meta-learning in NLP.

6 Experiments

We evaluate the models on few-shot generalization to new tasks and new domains of train tasks. Eval-

uation consist of a diverse set of NLP classification tasks from multiple domains: entity typing, sentiment classification, natural language inference and other text classification tasks. Our results¹ show that self-supervised meta-learning using SMLMT improves performance over self-supervised pre-training. Moreover, combining SMLMT with supervised tasks achieves the best generalization, improving over multi-task learning by up to 21%.

6.1 Implementation Details

SMLMT: We use the English Wikipedia dump, as of March 2019, to create SMLMT. This is similar to the dataset for pre-training of BERT (Devlin et al., 2018), which ensures that gains are not due to using more or diverse pre-training corpora (Liu et al., 2019b). The corpus is split into sentences and word-tokenized to create SMLMT. We run task creation offline and create about 2 Million SMLMT for meta-training, including a combination of 2, 3 and 4-way tasks. After task creation, the data is word-piece tokenized using the BERT-base cased model vocabulary for input to the models.

Supervised Tasks: Bansal et al. (2019) demonstrated that better feature learning from supervised tasks helps few-shot learning. Thus, we also evaluate multi-task learning and multi-task meta-learning for few-shot generalization. We also use GLUE tasks (Wang et al., 2018) and SNLI (Bowman et al., 2015) as the supervised tasks. Supervised tasks can be combined with SMLMT for meta-training (see 3). Note that since these are only a few supervised tasks (8 in this case) with a small label space, it is easy for meta-learning models to overfit to the supervised tasks (Yin et al., 2020) limiting generalization as we show in experiments.

Models: We evaluate the following models:

- (1) BERT: This is transformer model trained with self-supervised learning using MLM as the pre-training task on Wikipedia and BookCorpus. We use the cased base model (Devlin et al., 2018).
- (2) MT-BERT: This is a multi-task learning model trained on the supervised tasks. We follow Bansal et al. (2019) in training this model.
- (3) MT-BERT_{softmax}: This is the same model above where only the softmax layer is fine-tuned on downstream tasks.
- (4) LEOPARD: This is the meta-learning model proposed in Bansal et al. (2019) which is trained on only the supervised tasks.

¹Code and trained models: <https://github.com/iesl/metanlp>

(5) SMLMT: This is the meta-learning model (in 4) which is trained on the self-supervised SMLMT.
(6) Hybrid-SMLMT: This is the meta-learning model (in 4) trained on a combination of SMLMT and supervised tasks.

Note that all models share the same transformer architecture making the contribution from each component discernible. Moreover, SMLMT and Hybrid-SMLMT models use similar meta-learning algorithm as LEOPARD, so any improvements are due to the self-supervised meta-training. All model are initialized with pre-trained BERT for training.

Evaluation Methodology: We evaluate on few-shot generalization to multiple NLP tasks using the same set of tasks² considered in Bansal et al. (2019). Each target task consists of k examples per class, for $k \in \{4, 8, 16, 32\}$, and different tasks can have different number of classes. Since few-shot performance is sensitive to the few examples used in fine-tuning, each model is fine-tuned on 10 such k -shot support sets for a task, for each k , and the average performance with standard deviation is reported. Models are trained using their training procedures, without access to the target tasks, and are then fine-tuned for each of the k -shot task. Results for MT-BERT and LEOPARD are taken from Bansal et al. (2019).

Hyper-parameters: We follow the approach of Bansal et al. (2019) and use validation tasks for estimating hyper-parameters during fine-tuning for all baseline models. Note the meta-learning approach learn the learning rates during training and only require the number of epochs of fine-tuning to be estimated from the validation tasks. Detailed hyper-parameters are in Supplementary.

6.2 Results

6.2.1 Few-shot generalization to new tasks

We first evaluate performance on novel tasks not seen during training. The task datasets considered are: (1) entity typing: CoNLL-2003 (Sang and De Meulder, 2003), MIT-Restaurant (Liu et al., 2013); (2) rating classification (Bansal et al., 2019): 4 domains of classification tasks based on ratings from the Amazon Reviews dataset (Blitzer et al., 2007); (3) text classification: multiple social-media datasets from figure-eight³.

Results are presented in Table 1. Results on 2 domains of Rating are in Supplementary due to

space limitation. First, comparing models which don't use any supervised data, we see that on average across the 12 tasks, the meta-trained SMLMT performs better than BERT specially for small $k \in \{4, 8, 16\}$. Interestingly, the SMLMT model which doesn't use any supervised data, also outperforms even MT-BERT models which use supervised data for multi-task training. Next, comparing among all the models, we see that the Hybrid-SMLMT model performs best on average across tasks. For instance, on average 4-shot performance across tasks, Hybrid-SMLMT provides a relative gain in accuracy of 21.4% over the best performing MT-BERT baseline. Compared to LEOPARD, the Hybrid-SMLMT yields consistent improvements for all $k \in \{4, 8, 16, 32\}$ and demonstrates steady improvement in performance with increasing data (k). We note that on some tasks, such as Disaster, SMLMT is better than Hybrid-SMLMT. We suspect negative transfer from multi-task training on these tasks as also evidenced by the drop in performance of MT-BERT. These results show that SMLMT meta-training learns a better initial point that enables few-shot generalization.

6.2.2 Few-shot domain transfer

The tasks considered here had another domain of a similar task in the GLUE training tasks. Datasets used are (1) 4 domains of Amazon review sentiments (Blitzer et al., 2007), (2) Scitail, a scientific NLI dataset (Khot et al., 2018). Results on 2 domains of Amazon are in Supplementary due to space limitation. A relevant baseline here is MT-BERT_{reuse} which reuses the softmax layer from the related training task. This is a prominent approach to transfer learning with pre-trained models. Comparing Hybrid-SMLMT with variants of MT-BERT, we see that Hybrid-SMLMT performs comparable or better. Comparing with LEOPARD, we see that Hybrid-SMLMT generalizes better to new domains. LEOPARD performs worse than Hybrid-SMLMT on Scitail even though the supervised tasks are biased towards NLI, with 5 of the 8 tasks being variants of NLI tasks. This is due to meta-overfitting to the training domains in LEOPARD which is prevented through the regularization from SMLMT in Hybrid-SMLMT.

6.3 Analysis

Meta-overfitting: We study the extent of meta-overfitting in LEOPARD and Hybrid-SMLMT. Since these models learn the adaptation learning-

²Data: <https://github.com/iesl/leopard>

³<https://www.figure-eight.com/data-for-everyone/>

Task	N	k	BERT	SMLMT	MT-BERT _{softmax}	MT-BERT	LEOPARD	Hybrid-SMLMT
CoNLL	4	4	50.44 ± 08.57	46.81 ± 4.77	52.28 ± 4.06	55.63 ± 4.99	54.16 ± 6.32	57.60 ± 7.11
		8	50.06 ± 11.30	61.72 ± 3.11	65.34 ± 7.12	58.32 ± 3.77	67.38 ± 4.33	70.20 ± 3.00
		16	74.47 ± 03.10	75.82 ± 4.04	71.67 ± 3.03	71.29 ± 3.30	76.37 ± 3.08	80.61 ± 2.77
		32	83.27 ± 02.14	84.01 ± 1.73	73.09 ± 2.42	79.94 ± 2.45	83.61 ± 2.40	85.51 ± 1.73
MITR	8	4	49.37 ± 4.28	46.23 ± 3.90	45.52 ± 5.90	50.49 ± 4.40	49.84 ± 3.31	52.29 ± 4.32
		8	49.38 ± 7.76	61.15 ± 1.91	58.19 ± 2.65	58.01 ± 3.54	62.99 ± 3.28	65.21 ± 2.32
		16	69.24 ± 3.68	69.22 ± 2.78	66.09 ± 2.24	66.16 ± 3.46	70.44 ± 2.89	73.37 ± 1.88
		32	78.81 ± 1.95	78.82 ± 1.30	69.35 ± 0.98	76.39 ± 1.17	78.37 ± 1.97	79.96 ± 1.48
Airline	3	4	42.76 ± 13.50	42.83 ± 6.12	43.73 ± 7.86	46.29 ± 12.26	54.95 ± 11.81	56.46 ± 10.67
		8	38.00 ± 17.06	51.48 ± 7.35	52.39 ± 3.97	49.81 ± 10.86	61.44 ± 03.90	63.05 ± 8.25
		16	58.01 ± 08.23	58.42 ± 3.44	58.79 ± 2.97	57.25 ± 09.90	62.15 ± 05.56	69.33 ± 2.24
		32	63.70 ± 4.40	65.33 ± 3.83	61.06 ± 3.89	62.49 ± 4.48	67.44 ± 01.22	71.21 ± 3.28
Disaster	2	4	55.73 ± 10.29	62.26 ± 9.16	52.87 ± 6.16	50.61 ± 8.33	51.45 ± 4.25	55.26 ± 8.32
		8	56.31 ± 09.57	67.89 ± 6.83	56.08 ± 7.48	54.93 ± 7.88	55.96 ± 3.58	63.62 ± 6.84
		16	64.52 ± 08.93	72.86 ± 1.70	65.83 ± 4.19	60.70 ± 6.05	61.32 ± 2.83	70.56 ± 2.23
		32	73.60 ± 01.78	73.69 ± 2.32	67.13 ± 3.11	72.52 ± 2.28	63.77 ± 2.34	71.80 ± 1.85
Emotion	13	4	09.20 ± 3.22	09.84 ± 1.09	09.41 ± 2.10	09.84 ± 2.14	11.71 ± 2.16	11.90 ± 1.74
		8	08.21 ± 2.12	11.02 ± 1.02	11.61 ± 2.34	11.21 ± 2.11	12.90 ± 1.63	13.26 ± 1.01
		16	13.43 ± 2.51	12.05 ± 1.18	13.82 ± 2.02	12.75 ± 2.04	13.38 ± 2.20	15.17 ± 0.89
		32	16.66 ± 1.24	14.28 ± 1.11	13.81 ± 1.62	16.88 ± 1.80	14.81 ± 2.01	16.08 ± 1.16
Political Bias	2	4	54.57 ± 5.02	57.72 ± 5.72	54.32 ± 3.90	54.66 ± 3.74	60.49 ± 6.66	61.17 ± 4.91
		8	56.15 ± 3.75	63.02 ± 4.62	57.36 ± 4.32	54.79 ± 4.19	61.74 ± 6.73	64.10 ± 4.03
		16	60.96 ± 4.25	66.35 ± 2.84	59.24 ± 4.25	60.30 ± 3.26	65.08 ± 2.14	66.11 ± 2.04
		32	65.04 ± 2.32	67.73 ± 2.27	62.68 ± 3.21	64.99 ± 3.05	64.67 ± 3.41	67.30 ± 1.53
Political Audience	2	4	51.89 ± 1.72	57.94 ± 4.35	51.50 ± 2.72	51.47 ± 3.68	52.60 ± 3.51	57.40 ± 7.18
		8	52.80 ± 2.72	62.82 ± 4.50	53.53 ± 2.26	54.34 ± 2.88	54.31 ± 3.95	60.01 ± 4.54
		16	58.45 ± 4.98	64.57 ± 5.23	56.37 ± 2.19	55.14 ± 4.57	57.71 ± 3.52	63.11 ± 4.06
		32	55.31 ± 1.46	67.68 ± 3.12	53.09 ± 1.33	55.69 ± 1.88	52.50 ± 1.53	65.50 ± 3.78
Political Message	9	4	15.64 ± 2.73	16.16 ± 1.81	13.71 ± 1.10	14.49 ± 1.75	15.69 ± 1.57	16.74 ± 2.50
		8	13.38 ± 1.74	19.24 ± 2.32	14.33 ± 1.32	15.24 ± 2.81	18.02 ± 2.32	20.33 ± 1.22
		16	20.67 ± 3.89	21.91 ± 0.57	18.11 ± 1.48	19.20 ± 2.20	18.07 ± 2.41	22.93 ± 1.82
		32	24.60 ± 1.81	23.87 ± 1.72	18.67 ± 1.52	21.64 ± 1.78	19.87 ± 1.93	23.78 ± 0.54
Rating Electronics	3	4	39.27 ± 10.15	37.69 ± 4.82	39.89 ± 5.83	41.20 ± 10.69	51.71 ± 7.20	53.74 ± 10.17
		8	28.74 ± 08.22	39.98 ± 4.03	46.53 ± 5.44	45.41 ± 09.49	54.78 ± 6.48	56.64 ± 03.01
		16	45.48 ± 06.13	45.85 ± 4.72	48.71 ± 6.16	47.29 ± 10.55	58.69 ± 2.41	58.67 ± 03.73
		32	50.98 ± 5.89	50.86 ± 3.44	52.58 ± 2.48	53.49 ± 3.87	58.47 ± 5.11	61.42 ± 03.86
Rating Kitchen	3	4	34.76 ± 11.20	40.75 ± 7.33	40.41 ± 5.33	36.77 ± 10.62	50.21 ± 09.63	52.13 ± 10.18
		8	34.49 ± 08.72	43.04 ± 5.22	48.35 ± 7.87	47.98 ± 09.73	53.72 ± 10.31	58.13 ± 07.28
		16	47.94 ± 08.28	46.82 ± 3.94	52.94 ± 7.14	53.79 ± 09.47	57.00 ± 08.69	61.02 ± 05.55
		32	50.80 ± 04.52	51.71 ± 4.64	54.26 ± 6.37	53.23 ± 5.14	61.12 ± 04.83	64.69 ± 02.40
Overall Average		4	38.13	40.95	40.13	40.10	45.99	48.71
		8	36.99	46.37	45.89	44.25	50.86	53.70
		16	48.55	51.61	49.93	49.07	55.50	58.41
		32	55.30	56.23	52.65	55.42	57.02	60.81

Table 1: k -shot accuracy on novel tasks not seen in training. Models on left of separator don't use supervised data.

Task	k	BERT _{base}	SMLMT	MT-BERT _{softmax}	MT-BERT	MT-BERT _{reuse}	LEOPARD	Hybrid-SMLMT
Scitail	4	58.53 ± 09.74	50.68 ± 4.30	74.35 ± 5.86	63.97 ± 14.36	76.65 ± 2.45	69.50 ± 9.56	76.75 ± 3.36
	8	57.93 ± 10.70	55.60 ± 2.40	79.11 ± 3.11	68.24 ± 10.33	76.86 ± 2.09	75.00 ± 2.42	79.10 ± 1.14
	16	65.66 ± 06.82	56.51 ± 3.78	79.60 ± 2.31	75.35 ± 04.80	79.53 ± 2.17	77.03 ± 1.82	80.37 ± 1.44
	32	68.77 ± 6.27	62.38 ± 3.22	82.23 ± 1.12	74.87 ± 3.62	81.77 ± 1.13	79.44 ± 1.99	82.20 ± 1.34
Amazon Books	4	54.81 ± 3.75	55.68 ± 2.56	68.69 ± 5.21	64.93 ± 8.65	74.79 ± 6.91	82.54 ± 1.33	84.70 ± 0.42
	8	53.54 ± 5.17	60.23 ± 5.28	74.86 ± 2.17	67.38 ± 9.78	78.21 ± 3.49	83.03 ± 1.28	84.85 ± 0.52
	16	65.56 ± 4.12	62.92 ± 4.39	74.88 ± 4.34	69.65 ± 8.94	78.87 ± 3.32	83.33 ± 0.79	85.13 ± 0.66
	32	73.54 ± 3.44	71.49 ± 4.74	77.51 ± 1.14	78.91 ± 1.66	82.23 ± 1.10	83.55 ± 0.74	85.27 ± 0.36
Amazon DVD	4	54.98 ± 3.96	52.95 ± 2.51	63.68 ± 5.03	66.36 ± 7.46	71.74 ± 8.54	80.32 ± 1.02	83.28 ± 1.85
	8	55.63 ± 4.34	54.28 ± 4.20	67.54 ± 4.06	68.37 ± 6.51	75.36 ± 4.86	80.85 ± 1.23	83.91 ± 1.14
	16	58.69 ± 6.08	57.87 ± 2.69	70.21 ± 1.94	70.29 ± 7.40	76.20 ± 2.90	81.25 ± 1.41	83.71 ± 1.04
	32	66.21 ± 5.41	65.09 ± 4.37	70.19 ± 2.08	73.45 ± 4.37	79.17 ± 1.71	81.54 ± 1.33	84.15 ± 0.94

Table 2: k -shot domain transfer accuracy.

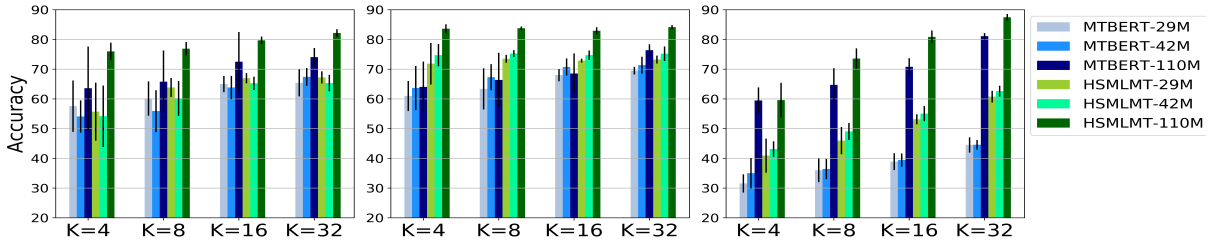


Figure 2: k -shot performance with number of parameters on Scitail (left), Amazon DVD (middle), and CoNLL (right). Larger models generalize better and Hybrid-SMLMT provides accuracy gains for all parameter sizes.

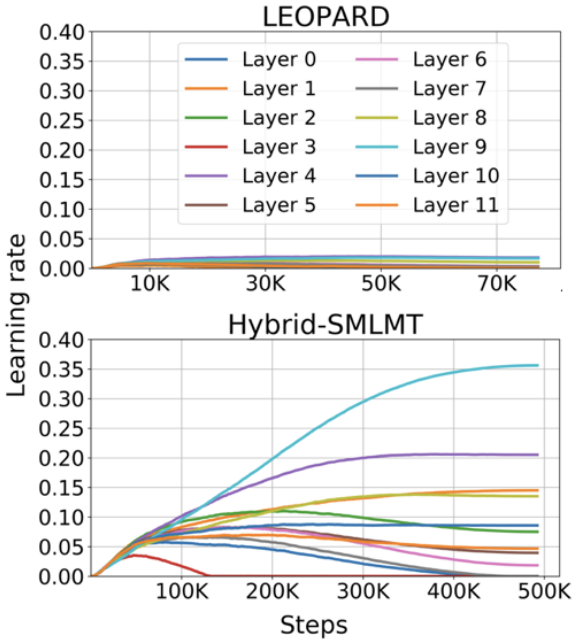


Figure 3: Learning rate trajectory during meta-training. LEOPARD learning-rates converge towards 0 for many layers, indicating meta-overfitting.

rates, we can study the learning rates trajectory during meta-training. Fig. 3 shows the results. We expect the learning rates to converge towards zero if the task-adaptation become irrelevant due to meta-overfitting. LEOPARD shows clear signs of meta-overfitting with much smaller learning rates which converge towards zero for most of the layers. Note that due to this, held-out validation during training is essential to enable any generalization (Bansal et al., 2019). Hybrid-SMLMT doesn’t show this phenomenon for most layers and learning rates converge towards large non-zero values even when we continue training for much longer. This indicates that SMLMT help in ameliorating meta-overfitting.

Effect of the number of parameters: We study how the size of the models affect few-shot performance. Recently, there has been increasing evidence that larger pre-trained models tend to gen-

eralize better (Devlin et al., 2018; Radford et al., 2019; Raffel et al., 2019). We explore whether this is true even in the few-shot regime. For this analysis we use the development data for 3 tasks: Scitail, Amazon DVD sentiment classification, and CoNLL entity typing. We consider the BERT base architecture with 110M parameters, and two smaller versions made available by Turc et al. (2019) consisting of about 29M and 42M parameters. We train versions of Hybrid-SMLMT as well as MTBERT corresponding to the smaller models. Results are presented in Fig. 2. Interestingly, we see that bigger models perform much better than the smaller models even when the target task had only 4 examples per class. Moreover, we see consistent and large performance gains from the meta-learned Hybrid-SMLMT, even for its smaller model variants. These results indicate that meta-training helps in data-efficient learning even with smaller models, and enables larger models to learn more generalizable representations.

Representation analysis: To probe how the representations in the proposed models are different from the representations in the self-supervised BERT model and multi-task BERT models, we performed CCA analysis on their representations (Raghu et al., 2017). We use the representations on the CoNLL and Scitail tasks for this analysis. Results on CoNLL task are in Fig. 4. First, we analyze the representation of the same model before and after fine-tuning on the target task. Interestingly, we see that the Hybrid-SMLMT model is closer to the initial point after task-specific fine-tuning than the BERT and MT-BERT models. Coupled with the better performance of Hybrid-SMLMT (in 6.2), this indicates a better initialization point for Hybrid-SMLMT. Note that the representations in lower layers are more similar before and after fine-tuning, and lesser in the top few layers. Next, we look at how representations differ across these mod-




Figure 4: CCA similarity for each transformer layer. Left: similarity before and after fine-tuning for the same model. Right: similarity between different pairs of models post fine-tuning. More results in Appendix.

els. We see that the models converge to different representations, where the lower layer representations are more similar and they diverge as we move towards the upper layers. In particular, note that this indicates that multi-task learning helps in learning different representations than self-supervised pre-training, and meta-learning model representations are different from the other models.

7 Conclusion

We introduced an approach to leverage unlabeled data to create meta-learning tasks for NLP. This enables better representation learning, learning key hyper-parameters like learning rates, demonstrates data-efficient fine-tuning, and ameliorates meta-overfitting when combined with supervised tasks. Through extensive experiments, we evaluated the proposed approach on few-shot generalization to novel tasks and domains and found that leveraging unlabelled data has significant benefits to enabling data-efficient generalization. This opens up the possibility of exploring large-scale meta-learning in NLP for various meta problems, including neural architecture search, continual learning, hyper-parameter learning, and more.

8 Acknowledgements

This work was supported in part by the Chan Zuckerberg Initiative, in part by the National Science Foundation under Grant No. IIS-1514053 and IIS-1763618, and in part by Microsoft Research Montréal. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2019. Learning to few-shot learn across diverse natural language classification tasks. *arXiv preprint arXiv:1911.03863*.
- Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, pages 6–8. Univ. of Texas.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pages 440–447.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Empirical Methods in Natural Language Processing*, pages 632–642.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(Feb):625–660.
- Chelsea Finn. 2018. *Learning to Learn with Gradients*. Ph.D. thesis, UC Berkeley.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135.

- Sebastian Flennerhag, Andrei A Rusu, Razvan Pascanu, Hujun Yin, and Raia Hadsell. 2019. Meta-learning with warped gradient descent. *arXiv preprint arXiv:1909.00025*.
- Jiatao Gu, Yong Wang, Yun Chen, Kyunghyun Cho, and Victor OK Li. 2018. Meta-learning for low-resource neural machine translation. *arXiv preprint arXiv:1808.08437*.
- Jiang Guo, Darsh Shah, and Regina Barzilay. 2018. Multi-source domain adaptation with mixture of experts. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4694–4703.
- 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 *Empirical Methods in Natural Language Processing*, pages 4803–4809.
- Nithin Holla, Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. Learning to learn to disambiguate: Meta-learning for few-shot word sense disambiguation. *arXiv preprint arXiv:2004.14355*.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339.
- Kyle Hsu, Sergey Levine, and Chelsea Finn. 2019. **Un-supervised learning via meta-learning**. In *International Conference on Learning Representations*.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. Scitail: A textual entailment dataset from science question answering. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. 2017. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*.
- Jingjing Liu, Panupong Pasupat, Scott Cyphers, and Jim Glass. 2013. Asgard: A portable architecture for multilingual dialogue systems. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8386–8390. IEEE.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*.
- 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. *arXiv preprint arXiv:1907.11692*.
- Luke Metz, Niru Maheswaranathan, Brian Cheung, and Jascha Sohl-Dickstein. 2019. **Learning unsupervised learning rules**. In *International Conference on Learning Representations*.
- Fei Mi, Minlie Huang, Jiyong Zhang, and Boi Faltings. 2019. Meta-learning for low-resource natural language generation in task-oriented dialogue systems. *arXiv preprint arXiv:1905.05644*.
- Tsendsuren Munkhdalai and Hong Yu. 2017. Meta networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2554–2563. JMLR. org.
- Alex Nichol and John Schulman. 2018. Reptile: a scalable metalearning algorithm. *arXiv preprint arXiv:1803.02999*, 2.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. 2017. Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *Advances in Neural Information Processing Systems*, pages 6076–6085.
- Sachin Ravi and Hugo Larochelle. 2017. Optimization as a model for few-shot learning. In *Proceedings of the International Conference on Learning Representations*.
- Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. 2018. Meta-learning with latent embedding optimization. *arXiv preprint arXiv:1807.05960*.
- Erik F Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.

Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850.

Jürgen Schmidhuber. 1987. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. thesis, Technische Universität München.

Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087.

Wilson L Taylor. 1953. “cloze procedure”: A new tool for measuring readability. *Journalism quarterly*, 30(4):415–433.

Sebastian Thrun and Lorien Pratt. 2012. *Learning to learn*. Springer Science & Business Media.

Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: The impact of student initialization on knowledge distillation. *ArXiv*, abs/1908.08962.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.

Lilian Weng. 2019. *Self-supervised representation learning*. lilianweng.github.io/lil-log.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.

Mingzhang Yin, George Tucker, Mingyuan Zhou, Sergey Levine, and Chelsea Finn. 2020. *Meta-learning without memorization*. In *International Conference on Learning Representations*.

Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomás Kociský, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. *Learning and evaluating general linguistic intelligence*. *CoRR*, abs/1901.11373.

Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauero, Haoyu Wang, and Bowen Zhou. 2018. Diverse few-shot text classification with multiple metrics. *arXiv preprint arXiv:1805.07513*.

A Appendix

A.1 Training Algorithm

The meta-training algorithm is given in 1. Note that π^w are the parameters for the warp layers and π are the remaining transformer parameters. $L_T(\cdot)$ is the cross-entropy loss for N -way classification in task T , calculated from the following prediction:

$$p(y|x) = \text{softmax} \{ \mathbf{W} h_\phi(f_\pi(x)) + \mathbf{b} \} \quad (5)$$

$g_\psi(\cdot)$ and h_ϕ are a two layer MLP with tanh non-linearity (Bansal et al., 2019).

Algorithm 1 Meta-Training

Require: SMLMT task distribution \mathcal{T} and supervised tasks \mathcal{S} , model parameters $\{\pi^w, \pi, \phi, \psi, \alpha\}$, adaptation steps G , learning-rate β , sampling ratio λ
Initialize θ with pre-trained BERT-base;

- 1: **while** not converged **do**
- 2: **for** task_batchsize times **do**
- 3: $t \sim \text{Bernoulli}(\lambda)$
- 4: $T \sim t \cdot \mathcal{T} + (1 - t) \cdot \mathcal{S}$
- 5: $\mathcal{D}^{tr} = \{(x_j, y_j)\} \sim T$
- 6: $C^n \leftarrow \{x_j | y_j = n\}; \quad N \leftarrow |C^n|$
- 7: $w^n, b^n \leftarrow \frac{1}{|C^n|} \sum_{x_j \in C^n} g_\psi(f_\pi(\mathcal{D}^{tr}))$
- 8: $\mathbf{W} \leftarrow [w^1; \dots; w^N]; \quad \mathbf{b} \leftarrow [b^1; \dots; b^N]$
- 9: $\theta \leftarrow \{\pi, \phi, \mathbf{W}, \mathbf{b}\}; \quad \theta^{(0)} \leftarrow \theta$
- 10: $\Theta \leftarrow \{\pi^w, \pi, \psi, \alpha\}$
- 11: $\mathcal{D}^{val} \sim T$
- 12: $q_T \leftarrow 0$
- 13: **for** $s := 0 \dots G - 1$ **do**
- 14: $\mathcal{D}_s^{tr} \sim T$
- 15: $\theta^{(s+1)} \leftarrow \theta^{(s)} - \alpha \nabla_{\theta} \mathcal{L}_T(\{\Theta, \theta^{(s)}\}, \mathcal{D}_s^{tr})$
- 16: $q_T \leftarrow q_T + \nabla_{\Theta} \mathcal{L}_T(\{\Theta, \theta^{(s+1)}\}, \mathcal{D}^{val})$
- 17: **end for**
- 18: **end for**
- 19: $\Theta \leftarrow \Theta - \beta \cdot \sum_T \frac{q_T}{G}$
- 20: **end while**

A.2 Additional Results

Table 5 shows the results the two additional domains of Rating classification, Table 6 shows the results for the two additional domains of Amazon sentiment classification. Fig.5 and Fig. 6 show the

Hyper-parameter	Value
Tasks per batch	4
Support samples per task	80
Query samples per task	10
Number of classes in SSLMT	[2,3,4]
d	256
Attention dropout	0.1
Hidden Layer Dropout	0.1
Outer Loop Learning Rate	1e-05
Adaptation Steps (G)	7
λ	0.5
Meta-training Epochs	1
Lowercase text	False
Sequence Length	128
Learning-rate Warmup	10% of steps

Table 3: Hyper-parameters.

CCA similarity on the two datasets: CoNLL and Scitail. Table 7 shows the accuracy for different model sizes on the three evaluation datasets: Scitail, Amazon DVD, CoNLL.

A.3 Datasets

Dataset splits and statistics are in Table 4.

Supervised Training Tasks: We selected the GLUE (Wang et al., 2018) benchmark tasks: MRPC, SST, MNLI (m/mm), QQP, QNLI, CoLA, RTE, and SNLI (Bowman et al., 2015) as the supervised training tasks for the meta-training phase. We used the standard train/dev/test split.

Test Tasks: These are same as the tasks used in Bansal et al. (2019).

A.4 Implementation Details

Training Hyper-parameters: Table 3 lists all the hyper-parameters for the Hybrid-SMLMT and SMLMT models. Both models use the same set of hyper-parameters, the difference being in the training tasks. Note, some hyper-parameters such as λ are not valid for SMLMT. We followed Devlin et al. (2018) in setting many hyper-parameters like dropouts, and Bansal et al. (2019) in setting hyper-parameters related to meta-learning. We use first-order MAML. Meta-training is run for only 1 epoch, so the model always trains on a new SMLMT in every batch. This corresponds to about 500,000 steps of updates during training.

Sampling for Hybrid-SMLMT: We restrict the word vocabulary for task creation with a term frequency of at least 50 in the corpus. This is then

Dataset	Labels	Train	Validation	Test
CoLA	2	8551	1042	—
MRPC	2	3669	409	—
QNLI	2	104744	5464	—
QQP	2	363847	40431	—
RTE	2	2491	278	—
SNLI	3	549368	9843	—
SST-2	2	67350	873	—
MNLI (m/mm)	3	392703	19649	—
Scitail	2	23,596	1,304	2,126
Amazon Sentiment Domains	2	800	200	1000
Airline	3	7320	—	7320
Disaster	2	4887	—	4887
Political Bias	2	2500	—	2500
Political Audience	2	2500	—	2500
Political Message	9	2500	—	2500
Emotion	13	20000	—	20000
CoNLL	4	23499	5942	5648
MIT-Restaurant	8	12474	—	2591

Table 4: Dataset statistics. Note that “—” indicates the corresponding split was not used.




Figure 5: Cross-model CCA similarity for each layer of the transformer after fine-tuning. Left plot is on CoNLL and right on Scitail.




Figure 6: CCA similarity for each layer of the same model before and after fine-tuning. Left plot is on CoNLL and right on Scitail.

used to create tasks in SMLMT as described. This word vocabulary is discarded at this point and the data is word-piece tokenized using the BERT-base cased model vocabulary for input to the models. Note that after a supervised task is selected to be sampled based on λ , it is sampled proportional to the square-root of the number of samples in the supervised tasks following Bansal et al. (2019).

Fine-tuning Hyper-parameter: We tune the number of fine-tuning epochs and batch-size using the development data of Scitail and Amazon Electronics tasks following (Bansal et al., 2019). Note

Task	N	k	BERT	SMLMT	MT-BERT _{softmax}	MT-BERT	LEOPARD	Hybrid-SMLMT
Rating Books	3	4	39.42 ± 07.22	34.96 ± 3.94	44.82 ± 9.00	38.97 ± 13.27	54.92 ± 6.18	57.80 ± 8.35
		8	39.55 ± 10.01	37.20 ± 4.15	51.14 ± 6.78	46.77 ± 14.12	59.16 ± 4.13	56.92 ± 5.64
		16	43.08 ± 11.78	43.62 ± 4.59	54.61 ± 6.79	51.68 ± 11.27	61.02 ± 4.19	63.33 ± 4.41
		32	52.21 ± 4.03	50.45 ± 3.28	54.97 ± 6.12	54.95 ± 4.82	64.11 ± 2.02	64.51 ± 3.06
Rating DVD	3	4	32.22 ± 08.72	38.26 ± 3.62	45.94 ± 7.48	41.23 ± 10.98	49.76 ± 9.80	52.08 ± 11.03
		8	36.35 ± 12.50	37.92 ± 3.61	46.23 ± 6.03	45.24 ± 9.76	53.28 ± 4.66	52.98 ± 07.84
		16	42.79 ± 10.18	41.87 ± 4.30	49.23 ± 6.68	45.19 ± 11.56	53.52 ± 4.77	56.70 ± 04.32
		32	48.61 ± 3.24	46.37 ± 4.91	51.16 ± 4.30	52.82 ± 3.41	55.49 ± 4.50	57.90 ± 03.93

Table 5: k -shot accuracy on novel tasks not seen in training. Results for 2 more rating tasks.

Task	k	BERT _{base}	SMLMT	MT-BERT _{softmax}	MT-BERT	MT-BERT _{reuse}	LEOPARD	Hybrid-SMLMT
Amazon Books	4	54.81 ± 3.75	55.68 ± 2.56	68.69 ± 5.21	64.93 ± 8.65	74.79 ± 6.91	82.54 ± 1.33	84.70 ± 0.42
	8	53.54 ± 5.17	60.23 ± 5.28	74.86 ± 2.17	67.38 ± 9.78	78.21 ± 3.49	83.03 ± 1.28	84.85 ± 0.52
	16	65.56 ± 4.12	62.92 ± 4.39	74.88 ± 4.34	69.65 ± 8.94	78.87 ± 3.32	83.33 ± 0.79	85.13 ± 0.66
	32	73.54 ± 3.44	71.49 ± 4.74	77.51 ± 1.14	78.91 ± 1.66	82.23 ± 1.10	83.55 ± 0.74	85.27 ± 0.36
Amazon Kitchen	4	56.93 ± 7.10	58.64 ± 4.68	63.07 ± 7.80	60.53 ± 9.25	75.40 ± 6.27	78.35 ± 18.36	80.70 ± 7.13
	8	57.13 ± 6.60	59.84 ± 3.66	68.38 ± 4.47	69.66 ± 8.05	75.13 ± 7.22	84.88 ± 01.12	84.74 ± 1.77
	16	68.88 ± 3.39	65.15 ± 5.83	75.17 ± 4.57	77.37 ± 6.74	80.88 ± 1.60	85.27 ± 01.31	85.32 ± 1.05
	32	78.71 ± 3.60	71.68 ± 4.34	76.64 ± 1.99	79.68 ± 4.10	82.18 ± 0.73	85.80 ± 0.70	86.33 ± 0.67
Amazon DVD	4	54.98 ± 3.96	52.95 ± 2.51	63.68 ± 5.03	66.36 ± 7.46	71.74 ± 8.54	80.32 ± 1.02	83.28 ± 1.85
	8	55.63 ± 4.34	54.28 ± 4.20	67.54 ± 4.06	68.37 ± 6.51	75.36 ± 4.86	80.85 ± 1.23	83.91 ± 1.14
	16	58.69 ± 6.08	57.87 ± 2.69	70.21 ± 1.94	70.29 ± 7.40	76.20 ± 2.90	81.25 ± 1.41	83.71 ± 1.04
	32	66.21 ± 5.41	65.09 ± 4.37	70.19 ± 2.08	73.45 ± 4.37	79.17 ± 1.71	81.54 ± 1.33	84.15 ± 0.94
Amazon Electronics	4	58.77 ± 6.10	56.40 ± 2.74	61.63 ± 7.30	64.13 ± 10.34	72.82 ± 6.34	74.88 ± 16.59	81.04 ± 1.77
	8	59.00 ± 5.78	62.06 ± 3.85	66.29 ± 5.36	64.21 ± 10.49	75.07 ± 3.40	81.29 ± 1.65	82.56 ± 0.81
	16	67.32 ± 4.18	64.57 ± 4.32	69.61 ± 3.54	71.12 ± 7.29	75.40 ± 2.43	81.86 ± 1.56	81.15 ± 2.39
	32	72.80 ± 4.30	70.10 ± 3.81	73.20 ± 2.14	72.30 ± 3.88	79.99 ± 1.58	82.40 ± 0.76	83.24 ± 1.14

Table 6: k -shot domain transfer accuracy for all 4 domains of Amazon sentiment classification.

	k	Small (29.1 M)		Medium (41.7 M)		Base (110.1 M)	
		MT-BERT	Our	MT-BERT	Our	MT-BERT	Our
Scitail	4	57.55 ± 8.64	55.70 ± 9.75	54.07 ± 5.43	54.17 ± 10.34	63.58 ± 14.04	75.98 ± 2.93
	8	60.13 ± 5.77	63.85 ± 3.19	55.88 ± 7.04	60.17 ± 5.86	65.77 ± 10.53	76.89 ± 2.28
	16	65.00 ± 2.73	66.98 ± 1.72	63.84 ± 3.91	65.23 ± 2.23	72.50 ± 10.01	79.71 ± 1.27
	32	65.40 ± 4.54	67.23 ± 2.05	67.40 ± 2.99	65.32 ± 2.76	74.04 ± 03.09	82.15 ± 1.29
Amazon DVD	4	60.99 ± 5.05	71.83 ± 6.69	63.66 ± 7.43	74.72 ± 3.74	64.04 ± 8.53	83.60 ± 1.49
	8	63.38 ± 6.91	73.49 ± 1.34	67.30 ± 4.39	75.24 ± 1.17	66.37 ± 9.12	83.75 ± 0.61
	16	67.99 ± 2.05	72.88 ± 0.66	70.73 ± 2.88	74.72 ± 1.58	68.52 ± 6.76	82.91 ± 1.20
	32	69.50 ± 1.28	73.24 ± 1.33	71.35 ± 2.83	75.20 ± 2.44	76.38 ± 2.00	84.13 ± 0.68
CoNLL	4	31.57 ± 3.06	40.91 ± 5.72	35.00 ± 5.11	43.12 ± 2.60	59.47 ± 4.40	59.60 ± 5.82
	8	35.97 ± 3.96	45.96 ± 4.58	36.40 ± 3.41	49.04 ± 2.84	64.72 ± 5.60	73.55 ± 3.44
	16	38.89 ± 2.84	53.14 ± 1.70	39.41 ± 2.21	55.05 ± 2.54	70.78 ± 2.92	80.85 ± 2.15
	32	44.50 ± 2.56	60.74 ± 1.96	44.57 ± 1.64	62.59 ± 1.83	81.09 ± 1.09	87.45 ± 1.12

Table 7: k -shot performance for three models sizes.

that best values are determined for each k . Epochs search range is [5, 10, 50, 100, 150, 200, 300, 400] and batch-size search range is [4, 8, 16]. The selected values, for $k = (4, 8, 16, 32)$, are: (1) Hybrid-SMLMT: epochs = (300, 350, 400, 200), batchsize = (8, 16, 8, 16); (2) SMLMT: epochs = (100, 200, 150, 200), batchsize = (8, 16, 8, 16). Expected overall average validation accuracy for these hyper-parameters, for $k \in (4, 8, 16, 32)$ are: (1) Hybrid-SMLMT: (0.80, 0.81, 0.83, 0.84); (2) SMLMT: (0.54, 0.56, 0.62, 0.68). Hyper-parameters for BERT, LEOPARD and MT-BERT

are taken from Bansal et al. (2019).

Training Hardware and Time: We train the SMLMT and Hybrid-SMLMT models on 4 V100 GPUs, each with 16GB memory. Owing to the warp layers, our training time per step and the GPU memory footprint is lower than LEOPARD (Bansal et al., 2019). However, our training typically runs much longer as the model doesn’t overfit unlike LEOPARD (see learning rate trajectory in main paper). Meta-training takes a total of 11 days and 14hours.