

SPoT: Better Frozen Model Adaptation through Soft Prompt Transfer

Tu Vu^{1,2}★ Brian Lester¹ Noah Constant¹ Rami Al-Rfou¹ Daniel Cer¹
Google Research¹
University of Massachusetts Amherst²
{ttvu, brianlester, nconstant, rmyeid, cer}@google.com
tuvu@cs.umass.edu

Abstract

As pre-trained language models have gotten larger, there has been growing interest in parameter-efficient methods to apply these models to downstream tasks. Building on the PROMPTTUNING approach of Lester et al. (2021), which learns task-specific soft prompts to condition a frozen language model to perform downstream tasks, we propose a novel prompt-based transfer learning approach called SPoT: **S**oft **P**rompt **T**ransfer. SPoT first learns a prompt on one or more source tasks and then uses it to initialize the prompt for a target task. We show that SPoT significantly boosts the performance of PROMPTTUNING across many tasks. More importantly, SPoT either matches or outperforms MODEL TUNING, which fine-tunes the entire model on each individual task, across all model sizes while being more parameter-efficient (up to 27,000× fewer task-specific parameters). We further conduct a large-scale study on task transferability with 26 NLP tasks and 160 combinations of source-target tasks, and demonstrate that tasks can often benefit each other via prompt transfer. Finally, we propose a simple yet efficient retrieval approach that interprets task prompts as task embeddings to identify the similarity between tasks and predict the most transferable source tasks for a given novel target task.

1 Introduction

The past few years have seen the rapid development of ever larger pre-trained language models, where it has repeatedly been shown that scaling up the model size is a key ingredient for achieving the best performance (Devlin et al., 2019; Raffel et al., 2020; Brown et al., 2020). While this trend has continued to push the boundaries of possibility across various NLP benchmarks, the sheer size of these models presents a major challenge for their

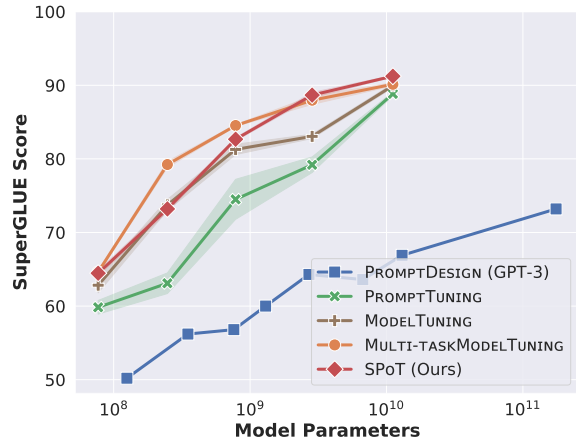


Figure 1: Our SPoT approach outperforms the vanilla PROMPTTUNING (Lester et al., 2021) and GPT-3’s few-shot PROMPTDESIGN (Brown et al., 2020) on the SuperGLUE benchmark by a substantial margin, obtaining competitive or significantly better results than MODEL TUNING across all model sizes. At the XXL model size, SPoT even outperforms the MULTI-TASK MODEL TUNING by over one point.

practical application. For 100B+ parameter models, fine-tuning and deploying a separate instance of the model for each downstream task would be prohibitively expensive.

To get around the infeasibility of fine-tuning, Brown et al. (2020) propose PROMPTDESIGN, where every downstream task is cast as a language modeling task and the *frozen* pre-trained model performs different tasks by conditioning on manual text prompts provided at inference time. Brown et al. (2020) demonstrate impressive few-shot performance with a single frozen GPT-3 model, although its performance depends highly on the choice of the prompt (Zhao et al., 2021) and still lags far behind state-of-the-art fine-tuning results.

More recent work has explored methods for learning soft prompts (Liu et al., 2021b; Qin and Eisner, 2021; Li and Liang, 2021; Lester et al., 2021), which can be seen as additional learnable parameters injected into the language model. Lester

★ Work done during an internship at Google Research.

et al. (2021) propose PROMPTTUNING, a simple method that learns a small task-specific prompt (a sequence of tunable tokens prepended to each example) for each downstream task during adaptation to condition the frozen language model to perform the task. Strikingly, as model capacity increases, PROMPTTUNING becomes competitive with MODEL TUNING, which fine-tunes the entire model on each downstream task. Nevertheless, at small and moderate model sizes (less than 11B parameters), there are still large gaps between PROMPTTUNING and MODEL TUNING.

In this paper, we propose SPoT: **S**oft **P**rompt **T**ransfer, a novel transfer learning approach in the context of prompt tuning. SPoT first trains a prompt on one or more source tasks and then uses the resulting prompt to initialize the prompt for a target (downstream) task. Our experiments show that SPoT offers significant improvements over PROMPTTUNING across tasks and model sizes. For instance, for the T5 BASE (220M parameter) and T5 XXL (11B parameter) models (Raffel et al., 2020), we obtain a +10.1 and +2.4 point average accuracy improvement respectively on the SuperGLUE benchmark (Wang et al., 2019b). More importantly, SPoT performs competitively or significantly better than MODEL TUNING across all model sizes (see Figure 1).

Motivated by these results, we investigate transferability between tasks, through the lens of task prompts. Our goal is to answer the following questions: (a) *For a given target task, when does initializing the prompt to that of a source task help improve performance?*; (b) *Can we use the task prompts to make more principled choices about which source tasks to use for a given novel target task?* To answer (a), we conduct a systematic study of the T5 model using 26 NLP tasks and 160 combinations of source and target tasks. Our results indicate that tasks can often benefit each other via prompt transfer. To address (b), we interpret the learned task prompts as *task embeddings* to construct a semantic space of tasks and formalize the similarity between tasks. We design an efficient retrieval algorithm that measures task embedding similarity, allowing practitioners to identify source tasks that are likely to yield positive transferability for a given novel target task.

To summarize, our contributions are as follows:

- We propose SPoT, a novel prompt-based transfer learning approach, and show that scale is not necessary for PROMPTTUNING to match the

performance of MODEL TUNING. SPoT yields competitive or significantly better results than MODEL TUNING across all model sizes.

- We conduct a large-scale and systematic study on task transferability, which demonstrates conditions under which tasks can benefit each other via prompt transfer.
- We propose an efficient retrieval approach that interprets task prompts as task embeddings to construct a semantic space of tasks, and measures task embedding similarity to identify which tasks could benefit each other.
- To facilitate future work on prompt-based learning, we will release our library of task prompts and pre-trained models, and provide practical recommendations for adapting our library to NLP practitioners.

2 Improving PROMPTTUNING with SPoT

To improve performance of PROMPTTUNING, SPoT introduces *source prompt tuning*, an intermediate training stage between language model pre-training and target prompt tuning (Figure 2, left), to learn a prompt on one or more source tasks (while still keeping the base model frozen), which is then used to initialize the prompt for a target task. Our approach retains all the computational benefits of PROMPTTUNING, i.e., for each target task, it only requires storing a small task-specific prompt while enabling the reuse of a single frozen pre-trained model for all tasks. In this section, we present a *task-agnostic* SPoT approach where a single transferred prompt is reused for all target tasks. In Section 3, we explore a *task-specific* approach that retrieves different prompts for different target tasks.

2.1 Experimental setup

Our frozen models are built on top of the pre-trained T5 checkpoints of all sizes: SMALL, BASE, LARGE, XL, XXL with 60M, 220M, 770M, 3B, and 11B parameters, respectively. In our experiments with SPoT, we leverage the LM adapted version of T5¹, which was found to be easier to optimize for PROMPTTUNING (Lester et al., 2021).

¹T5 1.1 checkpoints trained for an additional 100K steps using the “prefix LM” objective (Raffel et al., 2020), available at https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released_checkpoints.md




Figure 2: An illustration of our *task-agnostic* (left) and *task-specific* (right) SPoT approaches. **Left:** We learn a single prompt on one or more source tasks, which is then used to initialize the prompt for each target task. **Right:** We learn prompts for source tasks, and save early checkpoints as task embeddings and best checkpoints as source prompts. These form the keys and values of our prompt library. Given a novel target task, a user: (i) computes a task embedding, (ii) retrieves an optimal source prompt, and (iii) trains a target prompt, which is initialized with the source prompt.

2.1.1 Baselines

We compare SPoT to the following baselines:

PROMPTTUNING: The vanilla prompt tuning approach of Lester et al. (2021), where an independent prompt is directly trained on each target task.

MODEL TUNING & MULTI-TASK MODEL TUNING: We compare prompt tuning approaches to MODEL TUNING, the standard fine-tuning approach (Devlin et al., 2019; Raffel et al., 2020), where all of the pre-trained parameters are fine-tuned on each target task separately. For an apples-to-apples comparison, we also include MULTI-TASK MODEL TUNING, a more competitive baseline that first fine-tunes the entire model on the same mixture of source tasks used for SPoT before fine-tuning it individually on each target task.²

2.1.2 Evaluation datasets

We study downstream performance on a diverse set of tasks from the GLUE (Wang et al., 2019c) and SuperGLUE (Wang et al., 2019b) benchmarks (each with 8 datasets).³ Due to restricted test set

²In preliminary experiments, we found that using the original version of T5 1.1 (which is pre-trained exclusively on span corruption) results in better performance than using the LM adapted version for model tuning approaches. We therefore report results corresponding to the original T5 1.1 for MODEL TUNING and MULTI-TASK MODEL TUNING.

³These datasets include grammatical acceptability judgments (CoLA (Warstadt et al., 2019)), sentiment analysis (SST-2 (Socher et al., 2013)), paraphrasing/semantic similarity (MRPC (Dolan and Brockett, 2005), STS-B (Cer et al., 2017), QQP (Iyer et al., 2017)), natural language inference (MNLI (Williams et al., 2018), QNLI (Wang et al., 2019c)), RTE (Dagan et al., 2005, et seq.), CB (De Marneffe et al., 2019)), coreference resolution (WSC (Levesque et al., 2012)), sentence completion (COPA (Roemmele et al., 2011)), word sense disambiguation (WiC (Pilehvar and Camacho-Collados, 2019)), and question answering (MultiRC (Khashabi et al.,

access for GLUE and SuperGLUE, we train for a fixed number of steps and report results on the validation set associated with each dataset.⁴

2.1.3 Data for source prompt tuning

As with language model pre-training, the choice of training data is crucial for successful prompt transfer. To investigate the impact of source training data on downstream performance, we compare a diverse set of source tasks.

A single unsupervised learning task: We first consider training a prompt on a fraction of the C4 (Colossal Clean Crawled Corpus) dataset (Raffel et al., 2020) using the “prefix LM” objective discussed in Raffel et al. (2020). Although this task was used to pre-train our frozen T5 models already, it could still be helpful for learning a general-purpose prompt.

A single supervised learning task: Alternatively, we can train the prompt using a supervised task. We use either MNLI (Williams et al., 2018) or SQuAD (Rajpurkar et al., 2016) as single source tasks. MNLI was shown to be helpful for many sentence-level classification tasks (Phang et al., 2019), while SQuAD was found to generalize well to QA tasks (Talmor and Berant, 2019).

A multi-task mixture: So far, we have been training the prompt on a single source task. An alternative approach is multi-task training. Within T5’s unified text-to-text framework, this simply corresponds to mixing different datasets to-

2018), ReCoRD (Zhang et al., 2018), BoolQ (Clark et al., 2019)). We exclude the problematic WNLI (Levesque et al., 2012) dataset from GLUE, following Devlin et al. (2019).

⁴For tasks with multiple metrics, we use an average of the metrics.

gether. We explore mixing datasets from different NLP benchmarks or families of tasks, including GLUE, SuperGLUE, natural language inference (NLI), paraphrasing/semantic similarity, sentiment analysis, question answering on MRQA (Fisch et al., 2019), commonsense reasoning on RAINBOW (Lourie et al., 2021), machine translation, summarization, and natural language generation on GEM (Gehrmann et al., 2021).⁵ We create a mixture of source tasks from each of the NLP benchmarks/families of tasks above, using the examples-proportional mixing strategy in Raffel et al. (2020) with an artificial dataset size limit $\mathcal{K} = 2^{19}$ training examples. Finally, we include a mixture of C4 and all the labeled datasets in the NLP benchmark-/families of tasks mentioned above (55 datasets).

2.1.4 Training details

For both source and target prompt tuning, we closely follow the training procedure in Lester et al. (2021). Specifically, for each target task, the only new parameters introduced during tuning are a shared prompt $\rho \in \mathbb{R}^{\mathcal{L} \times \mathcal{E}}$ prepended to each (embedded) input sequence, where \mathcal{L} , \mathcal{E} are the prompt length and the embedding size, respectively. In all cases, we set $\mathcal{L} = 100$ tokens. We tune the prompt with a batch size of 32 for a fixed number of steps \mathcal{S} . We use the Adafactor optimizer (Shazeer and Stern, 2018) with default parameters except with a constant learning rate of 0.3, weight decay of $1e-5$, and parameter scaling turned off. The dropout probability is always kept at 0.1. All of our models are implemented using JAX (Bradbury et al., 2018) and Flax (Heek et al., 2020). During source prompt tuning, the prompt tokens are initialized using the SAMPLEDVOCAB scheme (where embeddings are sampled from the 5,000 most common tokens in T5’s vocabulary). During target prompt tuning, we initialize the prompts with the final prompt checkpoint from source prompt tuning. We save a checkpoint every 500 steps and report results on the checkpoint corresponding to the highest validation performance.

For PROMPTTUNING, following Lester et al. (2021), we initialize the prompt using the CLASS-LABEL scheme (where the prompt tokens are initialized with embeddings that represent an enumeration of the output classes) with a back off to the SAMPLEDVOCAB scheme to fill any remaining prompt positions). Training details for model tun-

⁵See Appendix A.1 for details about datasets.

Method	GLUE	SuperGLUE
BASELINE		
PROMPTTUNING	81.2 _{0.4}	66.6 _{0.2}
– longer tuning	78.4 _{1.7}	63.1 _{1.1}
SPoT with different source mixtures		
GLUE (8 tasks)	82.8 _{0.2}	73.2 _{0.3}
– longer tuning	82.0 _{0.2}	70.7 _{0.4}
C4	82.0 _{0.2}	67.7 _{0.3}
MNLI	82.5 _{0.0}	72.6 _{0.8}
SQuAD	82.2 _{0.1}	72.0 _{0.4}
SuperGLUE (8 tasks)	82.0 _{0.1}	66.6 _{0.2}
NLI (7 tasks)	82.6 _{0.1}	71.4 _{0.2}
Paraphrasing/similarity (4 tasks)	82.2 _{0.1}	69.7 _{0.5}
Sentiment (5 tasks)	81.1 _{0.2}	68.6 _{0.1}
MRQA (6 tasks)	81.8 _{0.2}	68.4 _{0.2}
RAINBOW (6 tasks)	80.3 _{0.6}	64.0 _{0.4}
Translation (3 tasks)	82.4 _{0.2}	65.3 _{0.1}
Summarization (9 tasks)	80.9 _{0.3}	67.1 _{1.0}
GEM (8 tasks)	81.9 _{0.2}	70.5 _{0.5}
All (C4 + 55 supervised tasks)	81.8 _{0.2}	67.9 _{0.9}

Table 1: GLUE and SuperGLUE results achieved by applying T5 BASE with different prompt tuning approaches. We report the mean and standard deviation (in the subscript) across three random seeds. SPoT significantly improves performance and stability of PROMPTTUNING across the two benchmarks.

ing can be found in Appendix A.

Longer tuning: While the number of tuning steps \mathcal{S} is set to 30K in Lester et al. (2021), we find that additional tuning is helpful when training on large datasets. As such, we set \mathcal{S} to $2^{18} = 262,144$, following Raffel et al. (2020), with the exception of ablation experiments (rows “– longer tuning”) in Table 1 which use $\mathcal{S} = 30K$.

2.2 Effect of SPoT

We compare the results of SPoT and other approaches in Figure 1 and Table 1. Below, we summarize and analyze each of our findings in detail.

SPoT significantly improves performance and stability of PROMPTTUNING: Our results on the GLUE and SuperGLUE benchmarks with T5 BASE are shown in Table 1. Overall, the results suggest that prompt transfer provides an effective means of improving performance for PROMPTTUNING. For example, the best-performing variant of SPoT outperforms the vanilla PROMPTTUNING approach on both GLUE and SuperGLUE by a substantial margin, obtaining +4.4 and +10.1 point average accuracy improvements, respectively. Our ablation study indicates that longer tuning is also an important ingredient for achieving our best performance, and

is complementary to prompt transfer. Additionally, when longer tuning is omitted, we observe that SPoT improves stability across runs.

Different source mixtures can lead to performance gains: Within our SPoT approach, we can compare the effectiveness of different source mixtures (see Table 1). Source prompt tuning on GLUE performs best on both GLUE and SuperGLUE, obtaining average scores of 82.8 and 73.2, respectively.⁶ Interestingly, unsupervised source prompt tuning on C4 (the same task used to pre-train our frozen models) still yields considerable improvements, even outperforming source prompt tuning on SuperGLUE for SuperGLUE tasks. Additionally, using MNLI or SQuAD as a single source dataset is particularly helpful for both GLUE and SuperGLUE. Finally, other source mixtures can also lead to significant gains, and some NLP benchmarks/families of tasks (e.g., NLI and paraphrasing/semantic similarity) are more beneficial than others.

SPoT helps close the gap with MODEL TUNING across all model sizes: We compare the performance of different approaches across model sizes on SuperGLUE in Figure 1. For SPoT, we show the performance resulting from source prompt training on a mixture of GLUE tasks. As shown in Lester et al. (2021), PROMPT TUNING becomes more competitive with scale, and at the XXL size (11B parameters), it even matches the performance of MODEL TUNING. However, at smaller model sizes, there are still large gaps between the two approaches. We show that SPoT helps close these gaps and even exceeds MODEL TUNING’s performance by a large margin at several model sizes, while retaining all the computational benefits conferred by PROMPT TUNING. Finally, SPoT produces competitive performance to the strong MULTI-TASK MODEL TUNING baseline while being more parameter-efficient in both multi-task source tuning and target tuning; at the XXL size, SPoT achieves the best average score of 91.2, +1.1 points better than MULTI-TASK MODEL TUNING, despite having 27,000× fewer task-specific parameters.

3 Investigating task transferability

Having established that prompt transfer is helpful for prompt tuning, we now shift our focus to investigating task transferability, through the lens of

⁶SuperGLUE tasks benefit less from source prompt tuning on SuperGLUE likely due to the small size of these datasets.

Name	Task type	Train
<i>16 source tasks</i>		
C4	language modeling	365M
DocNLI	NLI	942K
Yelp-2	sentiment analysis	560K
MNLI	NLI	393K
QQP	paraphrase detection	364K
QNLI	NLI	105K
ReCoRD	QA	101K
CxC	semantic similarity	88K
SQuAD	QA	88K
DROP	QA	77K
SST-2	sentiment analysis	67K
WinoGrande	commonsense reasoning	40K
HellaSWAG	commonsense reasoning	40K
MultiRC	QA	27K
CosmosQA	commonsense reasoning	25K
RACE	QA	25K
<i>10 target tasks</i>		
BoolQ	QA	9K
CoLA	grammatical acceptability	9K
STS-B	semantic similarity	6K
WiC	word sense disambiguation	5K
CR	sentiment analysis	4K
MRPC	paraphrase detection	4K
RTE	NLI	2K
WSC	coreference resolution	554
COPA	QA	400
CB	NLI	250

Table 2: Tasks used in our task transferability experiments, sorted by training dataset size.⁷

task prompts. To shed light on the transferability between different tasks, we conduct a large-scale empirical study with 26 NLP tasks (including one unsupervised task) and 160 combinations of source and target tasks. We demonstrate that tasks can help each other via prompt transfer in various situations, and task similarity plays an important role in determining transferability. Additionally, we show that by interpreting the task prompts as task embeddings, we can construct a semantic space of tasks and formulate a more rigorous notion of task similarity. Finally, we propose a retrieval algorithm that measures task embedding similarity to choose which source tasks to use for a given novel target task (Figure 2, right).

3.1 Experimental setup

We study a diverse set of 16 source datasets and 10 target datasets (see Table 2).⁸ We consider all

⁷C4 contains 365M documents but the actual number of examples created for the prefix LM task was much larger.

⁸In addition to the datasets mentioned in Section 2, we also use DocNLI (Yin et al., 2021), Yelp-2 (Zhang et al., 2015), CxC (Parekh et al., 2021), DROP (Dua et al., 2019), WinoGrande (Sakaguchi et al., 2020), HellaSWAG (Zellers et al., 2019), CosmosQA (Huang et al., 2019), RACE (Lai et al., 2017), and CR (Hu and Liu, 2004).

160 possible pairs of source and target datasets, and perform transfer from each source task to each target task.

3.1.1 Source and target tasks

The source tasks comprise one unsupervised task (C4) and 15 supervised tasks covering natural language inference (NLI), paraphrasing/semantic similarity, sentiment analysis, question answering (QA), and commonsense reasoning. All source tasks are data-rich or have been shown to yield positive transfer in prior work. To simulate a realistic scenario, we use low-resource tasks (less than 10K training examples) as target tasks. These tasks cover the above types of tasks, and additionally include grammatical acceptability, word sense disambiguation, and coreference resolution.

3.1.2 Training details

To limit computational costs, we use T5 BASE in all of our task transferability experiments. We perform 262,144 prompt tuning steps on each source task. The prompt checkpoint with the highest source task validation performance is selected to initialize prompts for different target tasks. Since the target datasets are small, we only perform 100K prompt tuning steps on each target task. We repeat each experiment three times with different random seeds. Other training details are the same as mentioned in Section 2.1.4.

3.1.3 Constructing a semantic space of tasks

Since only the prompt parameters are updated during prompt tuning on specific tasks, the task prompts likely encode task-specific knowledge. This suggests that they could be used to reason about the nature of tasks and their relationships. To test this idea, we interpret task prompts as *task embeddings* and construct a semantic space of tasks. Note that while we use the best prompt checkpoints from the source tasks for transfer to the target tasks, we use earlier prompt checkpoints as our task embeddings. This enables fast computation of task embeddings for novel target tasks. In our experiments, the task embedding is derived from a fixed prompt checkpoint, i.e., at 10K steps, for every task.⁹ We estimate the similarity between two tasks t^1, t^2 by measuring the similarity between their corresponding task embeddings e^1, e^2 , using the following metrics:

⁹Our experiments with other checkpoint alternatives yielded worse performance.

COSINE SIMILARITY OF AVERAGE TOKENS: We compute the cosine similarity between the average pooled representations of the prompt tokens: $sim(t^1, t^2) = \cos(\frac{1}{\mathcal{L}} \sum_i e_i^1, \frac{1}{\mathcal{L}} \sum_j e_j^2)$, where e_i^1, e_j^2 denote the respective prompt tokens of e^1, e^2 , and \cos denotes the cosine similarity.

PER-TOKEN AVERAGE COSINE SIMILARITY: We compute the average cosine similarity between every prompt token pair (e_i^1, e_j^2) : $sim(t^1, t^2) = \frac{1}{\mathcal{L}^2} \sum_i \sum_j \cos(e_i^1, e_j^2)$.

3.2 Predicting and exploiting transferability

We leverage our task embeddings to predict and exploit task transferability. Specifically, we explore methods to predict the most beneficial source tasks for a given target task and then make use of their prompts to improve performance on the target task. To enlarge our set of source prompts, we use the prompts from all the three different prompt tuning runs on each source task, resulting in 48 source prompts. Given a target task t with task embedding e^t , we rank all the source prompts ρ^s in descending order by the similarity between their corresponding task embeddings e^s and the target embedding e^t : $sim(e^s, e^t)$. We denote the ranked list of source prompts as ρ^{s_r} , where r denotes the rank ($r = 1, 2, \dots, 48$). We experiment with the following methods:

BEST OF TOP- k : We select the top- k source prompts and use each of them individually to initialize the target prompt. This procedure requires prompt tuning k times on the target task t , once for each source prompt. The best individual result is then used for evaluating the effectiveness of this method.

TOP- k WEIGHTED AVERAGE: We initialize the target prompt with a weighted average of the top- k source prompts $\sum_{r=1}^k \alpha_r \rho^{s_r}$ so that we only perform prompt tuning on the target task t once. The weights α_r are computed as $\alpha_r = \frac{sim(e^{s_r}, e^t)}{\sum_{l=1}^k sim(e^{s_l}, e^t)}$, where e^{s_r} denotes the corresponding task embedding of ρ^{s_r} .

TOP- k MULTI-TASK MIXTURE: We first identify the source tasks whose prompts are in the top- k prompts and mix their datasets and the target dataset together, using the examples-proportional mixing strategy of Raffel et al. (2020). Then, we

perform source prompt tuning on this multi-task mixture and use the final prompt checkpoint to initialize the prompt for target prompt tuning.

3.2.1 Evaluation

We report the average score across the target tasks achieved by using each of the methods described above. For each target task t , we measure the average and standard deviation of performance across the three different prompt tuning runs (which result in different task embeddings e^t). For comparison, we report the absolute and relative improvements over the baseline when prompt tuning on each target task from scratch (i.e., without any prompt transfer). Additionally, we include the oracle results achieved by using a brute-force search to identify the best possible out of 48 source prompts for each target task.

3.3 Effect of prompt-based task embeddings

In this section, we first analyze our task transferability results. Then, we demonstrate the effectiveness of using prompt-based task embeddings for representing tasks, and for predicting and exploiting task transferability.

Tasks can help each other via prompt transfer in various scenarios:

The results of our task transferability experiments (see Table 4 in Appendix C) indicate that in many cases, transferring the prompt from a source task to a target task (SOURCE \rightarrow TARGET) can provide significant gain on the target task. The transfer MNL1 \rightarrow CB yields the largest relative error reduction of 58.9% (from an average score of 92.7 to 97.0), followed by MNL1 \rightarrow COPA (29.1%) and ReCoRD \rightarrow WSC (20.0%). Using the best source prompt (out of 48) for each target task dramatically improves the average score across 10 target tasks from 74.7 to 80.7. Overall, our results show effective transfer from large source tasks that involve high-level reasoning about semantic relationships among sentences (e.g., MNL1), or when the source and target tasks are similar (e.g., CxC \rightarrow STS-B). Interestingly, positive transfer can occur in cases where the tasks are relatively dissimilar (e.g., ReCoRD \rightarrow WSC, SQuAD \rightarrow MRPC, CxC \rightarrow WiC).¹⁰

Task embeddings capture task relationships:

Figure 3 shows a hierarchically-clustered heatmap of cosine similarities between the task embeddings




Figure 3: A clustered heatmap of cosine similarities between the task embeddings of the 26 NLP tasks we study. Our prompt-based task embeddings capture task relationships: similar tasks are grouped together into clusters.




Figure 4: Correlation between task similarity and task transferability. Each point represents a source prompt. Cosine similarity (x-axis) is between that source task embedding and the indicated target task’s embedding (orange title), averaged over three runs for the target task. Relative error reduction (y-axis) measures the improvement on the target task when performing prompt transfer from that source prompt. We include the Pearson correlation coefficient (r) and p-value.

of the 26 NLP tasks we study using the COSINE SIMILARITY OF AVERAGE TOKENS metric.¹¹ We observe that our learned task embeddings capture many intuitive task relationships. Specifically, similar tasks

¹⁰Table 5 in Appendix C contains more cases.

¹¹To obtain the highest resolution of similarity between two tasks, we use the average of cosine similarities between their task embeddings obtained with all the three different prompt tuning runs (9 combinations).

are grouped together into clusters, including question answering (SQuAD, ReCoRD, and DROP; MultiRC and BoolQ), sentiment analysis (Yelp-2, SST-2, and CR), NLI (MNLI and CB; DocNLI and RTE), semantic similarity (STS-B and CxC), paraphrasing (MRPC and QQP), and commonsense reasoning (WinoGrande, HellaSWAG, and CosmosQA). We note that QNLI, which is an NLI task built from the SQuAD dataset, is not closely linked to SQuAD; this suggests that our task embeddings are more sensitive to the type of task than domain similarity. Interestingly, they also capture the unintuitive case of ReCoRD’s high transferability to WSC. Additionally, task embeddings that are derived from different prompts of the same task have high similarity scores (see Appendix D).

Correlation between task embedding similarity and task transferability:

Figure 4 shows how the relative error reduction on a target task changes as a function of the similarity between the source and target task embeddings. Overall, we find that there is a significantly positive correlation between task embedding similarity and task transferability on four (out of 10) target tasks we study, including STS-B ($p < 0.001$), CB ($p < 0.001$, not shown), WSC ($p < 0.01$), and RTE ($p < 0.05$), while it is less significant on the other tasks.

Task embeddings can be used to predict and exploit task transferability:

We compare the results of different methods for identifying which source prompts could be beneficial for a given target task in Table 3. We find that using the PER-TOKEN AVERAGE COSINE SIMILARITY metric yields better results than using the COSINE SIMILARITY OF AVERAGE TOKENS metric for small values of k (≤ 9). Our results also suggest that BEST OF TOP- k provides an effective means of predicting and exploiting task transferability. Simply choosing the source prompt whose associated task embedding has the highest similarity to the target embedding using the PER-TOKEN AVERAGE COSINE SIMILARITY metric improves over the baseline by a large margin (from an average score of 74.7 to 76.7, a 12.1% average relative error reduction). Trying all the top-3 (out of 48) source prompts for each target task yields an average score of 77.5. With larger values of k , we can retain most of the benefits of oracle selection of source prompts (80% of the gain in terms of average score with $k = 9$ and 90% with $k = 15$), while still eliminating over 2/3 of the candidate

Method	Change		Avg. score
	Abs.	Rel.	
BASELINE	-	-	74.7 _{0.7}
BRUTE-FORCE SEARCH ($k = 48$)			
ORACLE	6.0 _{0.5}	26.5 _{1.1}	80.7 _{0.0}
COSINE SIMILARITY OF AVERAGE TOKENS			
BEST OF TOP- k			
$k = 1$	1.5 _{0.5}	11.7 _{1.1}	76.2 _{0.1}
$k = 3$	2.7 _{0.6}	16.6 _{1.1}	77.4 _{0.3}
$k = 6$	3.8 _{0.1}	20.0 _{1.1}	78.5 _{0.5}
$k = 9$	4.5 _{0.4}	22.2 _{1.1}	79.2 _{0.1}
$k = 12$	5.0 _{0.9}	23.6 _{2.2}	79.7 _{0.4}
$k = 15$	5.4 _{0.8}	24.9 _{1.8}	80.1 _{0.3}
PER-TOKEN AVERAGE COSINE SIMILARITY			
BEST OF TOP- k			
$k = 1$	2.0 _{0.4}	12.1 _{1.1}	76.7 _{0.7}
$k = 3$	2.9 _{0.6}	17.0 _{0.6}	77.5 _{0.4}
$k = 6$	4.5 _{0.5}	22.1 _{1.2}	79.2 _{0.1}
$k = 9$	4.6 _{0.5}	22.6 _{0.9}	79.5 _{0.2}
$k = 12$	5.0 _{0.6}	23.5 _{1.4}	79.6 _{0.1}
$k = 15$	5.3 _{0.9}	24.5 _{2.2}	80.0 _{0.4}
TOP- k WEIGHTED AVERAGE			
best $k = 3$	1.9 _{0.5}	11.5 _{2.7}	76.6 _{0.1}
TOP- k MULTI-TASK MIXTURE			
best $k = 12$	3.1 _{0.5}	15.3 _{2.8}	77.8 _{0.1}

Table 3: Task embeddings provides an effective means of predicting and exploiting task transferability. Using BEST OF TOP- k with $k = 3$ improves over the baseline by +2.8 points. With larger values of k (≤ 15), we can retain most of the benefits conferred by prompt transfer. For TOP- k WEIGHTED AVERAGE and TOP- k MULTI-TASK MIXTURE, we experiment with different values of $k \in \{3, 6, 9, 12\}$ and report the best results.

source prompts. Although this approach requires prompt tuning on the target task k times, the cost of prompt tuning is relatively inexpensive, compared to model tuning. TOP- k WEIGHTED AVERAGE has similar average performance to BEST OF TOP- k with $k = 1$, but achieves lower variance. Thus, this may be an appealing alternative to BEST OF TOP- k in scenarios where multiple tuning runs on the target task are prohibited. Finally, TOP- k MULTI-TASK MIXTURE also provides a means of obtaining strong performance with an average score of 77.8, even outperforming BEST OF TOP- k with $k \leq 3$.

4 Related Work

Parameter-efficient transfer learning & language model prompting

Pre-trained language models have been shown to be an effective means for improving state-of-the-art results on many NLP benchmarks (Devlin et al., 2019; Liu et al., 2019b; Yang et al., 2019; Lan et al., 2020; Raffel et al., 2020; Brown et al., 2020; He et al., 2021). However, MODEL TUNING (*a.k.a* fine-tuning)—the cur-

rent dominant approach for applying these models to downstream tasks—can become impractical, as fine-tuning all of the pre-trained parameters for each task can be prohibitively expensive, especially as model size continues to increase.

To address this issue, early work uses compression techniques, such as knowledge distillation (Sanh et al., 2019; Jiao et al., 2020; Sun et al., 2020) and model pruning (Fan et al., 2020; Sanh et al., 2020; Chen et al., 2020), to obtain lightweight pre-trained models. Other work involves updating only small parts of the language model (Zaken et al., 2021) or training task-specific modules, such as adapters (Houlsby et al., 2019; Karimi Mahabadi et al., 2021) and/or low-rank structures (Mahabadi et al., 2021; Hu et al., 2021), while keeping most or all of the pre-trained parameters fixed. Notably, Brown et al. (2020) demonstrate remarkable few-shot learning performance with a single frozen GPT-3 model using PROMPTDESIGN, where every task is cast as feeding the model a manual text prompt at inference time for context and asking it to produce some output text.

Several efforts have since focused on developing prompt-based learning approaches with carefully handcrafted prompts (Schick and Schütze, 2021), prompt mining and paraphrasing (Jiang et al., 2020b), gradient-based search for improved prompts (Shin et al., 2020), and automatic prompt generation (Gao et al., 2021). The use of hard prompts, however, was found to be sub-optimal and sensitive, i.e., there is no obvious correlation between downstream performance and the prompt format, and minor changes in the prompt can lead to significant differences in downstream performance (Liu et al., 2021b). As such, recent work has shifted toward learning soft prompts (Liu et al., 2021b; Qin and Eisner, 2021; Li and Liang, 2021; Lester et al., 2021), which can be seen as some additional learnable parameters injected into the language model. We refer readers to Liu et al. (2021a) for a recent survey on prompt-based learning research.

Concurrent work (Gu et al., 2021) also explores the effectiveness of prompt pre-training. Their approach uses hand-crafted pre-training tasks tailored to different types of downstream tasks, which limits its application to novel downstream tasks. In contrast, we use existing tasks as source tasks and show that prompt transfer can confer benefits even when there are mismatches (e.g., task type, input/output

format) between the source and target tasks. Their work also focuses on the few-shot setting, whereas we work in context of larger datasets. Additionally, we study task transferability and demonstrate that tasks can often help each other via prompt transfer, and task prompts can be interpreted as task embeddings to formalize task similarity to identify which tasks could benefit each other.

Task transferability We also build on existing work on task transferability in NLP (Phang et al., 2019; Wang et al., 2019a; Liu et al., 2019a; Talmor and Berant, 2019; Pruksachatkun et al., 2020; Vu et al., 2020; Poth et al., 2021) and computer vision (Zamir et al., 2018; Achille et al., 2019; Yan et al., 2020). Prior work shows effective transfer from data-rich source tasks (Phang et al., 2019), those that require complex reasoning and inference (Pruksachatkun et al., 2020), or those that are similar to the target task (Vu et al., 2020). There have also been efforts to predict transferability between tasks (Bingel and Søgaard, 2017; Vu et al., 2020; Poth et al., 2021). Vu et al. (2020) use task embeddings derived from either the input text or the diagonal Fisher information matrix of the language model, while Poth et al. (2021) explore adapter-based approaches. Here, our use of T5 allows us to better model the space of tasks, as every task is cast into a unified text-to-text format and the same model (without task-specific components) is used across tasks. Additionally, prompt-based task embeddings are comparatively cheaper to obtain.

5 Conclusion

In this paper, we study transfer learning in the context of prompt tuning. We show that scale is not necessary for PROMPTTUNING to match the performance of MODEL TUNING. Our SPOT approach matches or even exceeds the performance of MODEL TUNING by a large margin across model sizes while being more parameter-efficient (up to $27,000\times$ fewer task-specific parameters). Our large-scale study on task transferability indicates that tasks can benefit each other via prompt transfer in various scenarios. Finally, we demonstrate that task prompts can be interpreted as task embeddings to formalize the similarity between tasks. We propose a simple yet efficient retrieval approach that measures task similarity to identify which source tasks could confer benefits to a novel target task. Taken as a whole, we hope that our work will spur more research into prompt-based transfer learning.

Acknowledgements

We thank Mohit Iyyer, Sebastian Ruder, Kalpesh Krishna, Thang Luong, Quoc Le, and the members of the Descartes team for helpful discussion and feedback. We would also like to thank Lucas Dixon, Grady Simon, and Slav Petrov for their comments on this manuscript. Finally, we are grateful to Vamsi Aribandi for his work on preprocessing several datasets used in our experiments.

References

- Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charles C. Fowlkes, Stefano Soatto, and Pietro Perona. 2019. [Task2vec: Task embedding for meta-learning](#). In *Proceedings of the IEEE International Conference on Computer Vision (ICCV 2019)*, pages 6430–6439.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*.
- Joachim Bingel and Anders Søgaard. 2017. [Identifying beneficial task relations for multi-task learning in deep neural networks](#). In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*, pages 164–169.
- Yonatan Bisk, Rowan Zellers, Ronan Le bras, Jianfeng Gao, and Yejin Choi. 2020. [Piqa: Reasoning about physical commonsense in natural language](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2020)*, 34(05):7432–7439.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation (WMT 2014)*, pages 12–58.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yebes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation (WMT 2016)*, pages 131–198.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Barry Haddow, Matthias Huck, Chris Hokamp, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Carolina Scarton, Lucia Specia, and Marco Turchi. 2015. [Findings of the 2015 workshop on statistical machine translation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation (WMT 2015)*, pages 1–46.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 632–642.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. [JAX: composable transformations of Python+NumPy programs](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, volume 33, pages 1877–1901.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval 2017)*, pages 1–14.
- Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020. [The lottery ticket hypothesis for pre-trained bert networks](#). In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, volume 33, pages 15834–15846.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (ACL 2019)*, pages 2924–2936.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The pascal recognising textual entailment challenge](#). In *Proceedings of the 1st International*

- Conference on Machine Learning Challenges: Evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment (MLCW 2005)*, page 177–190.
- Marie-Catherine De Marneffe, Mandy Simons, and Judith Tonhauser. 2019. The commitmentbank: Investigating projection in naturally occurring discourse. In *Proceedings of Sinn und Bedeutung 23 (SuB 2018)*, volume 23, pages 107–124.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. **GoEmotions: A dataset of fine-grained emotions**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 4040–4054.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 4171–4186.
- William B. Dolan and Chris Brockett. 2005. **Automatically constructing a corpus of sentential paraphrases**. In *Proceedings of the Third International Workshop on Paraphrasing (IWP 2005)*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. **DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs**. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 2368–2378.
- Matthew Dunn, Levent Sagun, Mike Higgins, V Ugur Guney, Volkan Cirik, and Kyunghyun Cho. 2017. **Searchqa: A new q&a dataset augmented with context from a search engine**. *arXiv preprint arXiv:1704.05179*.
- Ondřej Dušek, David M. Howcroft, and Verena Rieser. 2019. **Semantic noise matters for neural natural language generation**. In *Proceedings of the 12th International Conference on Natural Language Generation (INLG 2019)*, pages 421–426.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. **Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*, pages 1074–1084.
- Angela Fan, Edouard Grave, and Armand Joulin. 2020. **Reducing transformer depth on demand with structured dropout**. In *Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*.
- Adam Fisch, Alon Talmor, Robin Jia, Minjoon Seo, Eunsol Choi, and Danqi Chen. 2019. **MRQA 2019 shared task: Evaluating generalization in reading comprehension**. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering (MRQA 2019)*, pages 1–13.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. **Making pre-trained language models better few-shot learners**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL 2021)*, pages 3816–3830.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. **Creating training corpora for NLG micro-planners**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 179–188.
- Sebastian Gehrmann, Tosin Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Anuluwapo Aremu, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh Dhole, Wanyu Du, Esin Durmus, Ondřej Dušek, Chris Chinenye Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Mihir Kale, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Andre Niyongabo Rubungo, Salomey Osei, Ankur Parikh, Laura Perez-Beltrachini, Niranjana Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. **The GEM benchmark: Natural language generation, its evaluation and metrics**. In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 96–120.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. **SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization**. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization (NewSum 2019)*, pages 70–79.
- Alec Go, Richa Bhayani, and Lei Huang. 2009. **Twitter sentiment classification using distant supervision**. *CS224N Project Report, Stanford*.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.

- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018)*, pages 708–719.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. [PPT: Pre-trained prompt tuning for few-shot learning](#). *arXiv preprint arXiv:2109.04332*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *Proceedings of the 9th International Conference on Learning Representations (ICLR 2021)*.
- Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas Steiner, and Marc van Zee. 2020. [Flax: A neural network library and ecosystem for JAX](#).
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Proceedings of the 29th Conference on Neural Information Processing Systems (NeurIPS 2020)*, volume 28.
- 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 *Proceedings of the 36th International Conference on Machine Learning (PMLR 2019)*, volume 97, pages 2790–2799.
- 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](#). *arXiv preprint arXiv:2106.09685*.
- Minqing Hu and Bing Liu. 2004. [Mining and summarizing customer reviews](#). In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2004)*, page 168–177.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019)*, pages 2391–2401.
- Shankar Iyer, Nikhil Dandekar, and Kornél Csernai. 2017. [First Quora Dataset Release: Question pairs](#).
- Chao Jiang, Mounica Maddela, Wuwei Lan, Yang Zhong, and Wei Xu. 2020a. [Neural CRF model for sentence alignment in text simplification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 7943–7960.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020b. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics (TACL 2020)*, 8:423–438.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. [TinyBERT: Distilling BERT for natural language understanding](#). In *Findings of the Association for Computational Linguistics (Findings of EMNLP 2020)*, pages 4163–4174.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 1601–1611.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021. [Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*, pages 565–576.
- Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. [Looking beyond the surface: A challenge set for reading comprehension over multiple sentences](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018)*, pages 252–262.
- Anastassia Kornilova and Vladimir Eidelman. 2019. [BillSum: A corpus for automatic summarization of US legislation](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization (NewSum 2019)*, pages 48–56.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics (TACL 2019)*, 7:452–466.
- Faisal Ladhak, Esin Durmus, Claire Cardie, and Kathleen McKeown. 2020. [WikiLingua: A new benchmark dataset for cross-lingual abstractive summarization](#). In *Findings of the Association for Computational Linguistics (Findings of EMNLP 2020)*, pages 4034–4048.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale Reading comprehension dataset from examinations](#). In

- Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, pages 785–794.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. **ALBERT: A lite BERT for self-supervised learning of language representations**. In *Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. **The power of scale for parameter-efficient prompt tuning**. *arXiv preprint arXiv:2104.08691*.
- Hector J. Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)*, page 552–561.
- Xiang Lisa Li and Percy Liang. 2021. **Prefix-tuning: Optimizing continuous prompts for generation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL 2021)*, pages 4582–4597.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. **CommonGen: A constrained text generation challenge for generative commonsense reasoning**. In *Findings of the Association for Computational Linguistics (Findings of EMNLP 2020)*, pages 1823–1840.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019a. Linguistic knowledge and transferability of contextual representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 1073–1094.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. **Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing**. *arXiv preprint arXiv:2107.13586*.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. **Gpt understands, too**. *arXiv preprint arXiv:2103.10385*.
- 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*.
- Nicholas Lourie, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. **Unicorn on rainbow: A universal commonsense reasoning model on a new multitask benchmark**. *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2021)*, 35(15):13480–13488.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021. **Compacter: Efficient low-rank hypercomplex adapter layers**. *arXiv preprint arXiv:2106.04647*.
- Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. 2021. **DART: Open-domain structured data record to text generation**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2021)*, pages 432–447.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. **Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*, pages 1797–1807.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. **Adversarial NLI: A new benchmark for natural language understanding**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 4885–4901.
- Zarana Parekh, Jason Baldridge, Daniel Cer, Austin Waters, and Yinfei Yang. 2021. **Crisscrossed captions: Extended intramodal and intermodal semantic similarity judgments for MS-COCO**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021)*, pages 2855–2870.
- Jason Phang, Thibault Févry, and Samuel R Bowman. 2019. **Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks**. *arXiv preprint arXiv:1811.01088*.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. **WiC: the word-in-context dataset for evaluating context-sensitive meaning representations**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2019)*, pages 1267–1273.
- Clifton Poth, Jonas Pfeiffer, Andreas Rücklé, and Iryna Gurevych. 2021. **What to pre-train on? efficient intermediate task selection**. *arXiv preprint arXiv:2104.08247*.
- Yada Pruksachatkun, Jason Phang, Haokun Liu, Phu Mon Htut, Xiaoyi Zhang, Richard Yuanzhe Pang, Clara Vania, Katharina Kann, and Samuel R. Bowman. 2020. **Intermediate-task transfer learning with pretrained language models: When and why**

- does it work? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 5231–5247.
- Guanghui Qin and Jason Eisner. 2021. [Learning how to ask: Querying LMs with mixtures of soft prompts](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2021)*, pages 5203–5212.
- 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](#). *Journal of Machine Learning Research (JMLR 2020)*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*, pages 2383–2392.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2020)*, 34(05):8689–8696.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. [Choice of plausible alternatives: An evaluation of commonsense causal reasoning](#). In *Proceedings of the 25th AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning (AAAI Spring Symposium 2011)*.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. [A neural attention model for abstractive sentence summarization](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*, pages 379–389.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. [Winogrande: An adversarial winograd schema challenge at scale](#). *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2020)*, 34(05):8732–8740.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#). In *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing (EMC²2019)*.
- Victor Sanh, Thomas Wolf, and Alexander Rush. 2020. [Movement pruning: Adaptive sparsity by fine-tuning](#). In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS 2020)*, volume 33, pages 20378–20389.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP 2019)*, pages 4463–4473.
- Timo Schick and Hinrich Schütze. 2021. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2021)*, pages 2339–2352.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017)*, pages 1073–1083.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). *arXiv preprint arXiv:1804.04235*.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 4222–4235.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*, pages 1631–1642.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. [MobileBERT: a compact task-agnostic BERT for resource-limited devices](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL 2020)*, pages 2158–2170.
- Alon Talmor and Jonathan Berant. 2019. [MultiQA: An empirical investigation of generalization and transfer in reading comprehension](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL 2019)*, pages 4911–4921.
- Adam Trischler, Tong Wang, Xingdi Yuan, Justin Harris, Alessandro Sordani, Philip Bachman, and Kaheer Suleman. 2017. [NewsQA: A machine comprehension dataset](#). In *Proceedings of the Workshop on Representation Learning for NLP (RepL4NLP 2017)*, pages 191–200.
- Tu Vu, Tong Wang, Tsendsuren Munkhdalai, Alessandro Sordani, Adam Trischler, Andrew Mattarella-Micke, Subhansu Maji, and Mohit Iyyer. 2020. [Exploring and predicting transferability across NLP](#)

- tasks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 7882–7926.
- Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019a. Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL 2019)*, pages 4465–4476.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019b. Superglue: A stickier benchmark for general-purpose language understanding systems. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS 2019)*, volume 32, pages 3266–3280.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2019c. Glue: A multi-task benchmark and analysis platform for natural language understanding. *Proceedings of the 7th International Conference on Learning Representations (ICLR 2019)*.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics (TACL 2019)*, 7:625–641.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2018)*, pages 1112–1122.
- Xi Yan, David Acuna, and Sanja Fidler. 2020. Neural data server: A large-scale search engine for transfer learning data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2020)*, pages 3893–3902.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Proceedings of the 33th Conference on Neural Information Processing Systems (NeurIPS 2019)*, volume 32, pages 5753–5763.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2018)*, pages 2369–2380.
- Wenpeng Yin, Dragomir Radev, and Caiming Xiong. 2021. DocNLI: A large-scale dataset for document-level natural language inference. In *Findings of the Association for Computational Linguistics (Findings of ACL-IJCNLP 2021)*, pages 4913–4922.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*.
- Amir R. Zamir, Alexander Sax, William Shen, Leonidas J. Guibas, Jitendra Malik, and Silvio Savarese. 2018. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2018)*, pages 3712–3722.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*, pages 4791–4800.
- Rui Zhang and Joel Tetreault. 2019. This email could save your life: Introducing the task of email subject line generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*, pages 446–456.
- Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. 2018. Record: Bridging the gap between human and machine commonsense reading comprehension. *arXiv preprint arXiv:1810.12885*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of the 29th Conference on Neural Information Processing Systems (NeurIPS 2015)*, volume 28, pages 649–657.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, volume 139 of PMLR, pages 12697–12706.

Appendices

A Additional training details

For model tuning approaches, we use the default hyperparameters for T5 (Raffel et al., 2020), i.e., learning rate 0.001, Adafactor optimizer with pre-training parameter states restored, and dropout probability 0.1. To improve the model tuning baselines, we perform a sweep over the batch size hyperparameter and select 2^{16} tokens per batch, following Lester et al. (2021).

B Source datasets used in our SPoT experiments in Section 2

Figure 5 displays the datasets used in our PROMPT-TUNING with SPoT experiments in Section 2. In addition to the C4 unlabeled dataset (Raffel et al., 2020), we use 55 labeled datasets. These datasets come from common NLP benchmarks/families of tasks, namely:

- GLUE (Wang et al., 2019c), including CoLA (Warstadt et al., 2019), SST-2 (Socher et al., 2013), MRPC (Dolan and Brockett, 2005), QQP (Iyer et al., 2017), STS-B (Cer et al., 2017), MNLI (Williams et al., 2018), QNLI (Wang et al., 2019c), and RTE (Dagan et al., 2005, et seq.).
- SuperGLUE (Wang et al., 2019b), including BoolQ (Clark et al., 2019), CB (De Marneffe et al., 2019), COPA (Roemmele et al., 2011), MultiRC (Khashabi et al., 2018), ReCoRD (Zhang et al., 2018), RTE, WiC (Pilehvar and Camacho-Collados, 2019), and WSC (Levesque et al., 2012).
- Natural language inference (NLI), including ANLI (Nie et al., 2020), CB, DocNLI (Yin et al., 2021), MNLI, QNLI, RTE, and SNLI (Bowman et al., 2015).
- Paraphrasing/semantic similarity, including CxC (Parekh et al., 2021), MRPC, QQP, and STS-B.
- Sentiment analysis, including CR (Hu and Liu, 2004), Goemotions (Demszky et al., 2020), Sentiment140 (Go et al., 2009), SST-2, and Yelp-2 (Zhang et al., 2015).
- Question answering on MRQA (Fisch et al., 2019), including SQuAD (Rajpurkar et al.,

2016), NewsQA (Trischler et al., 2017), TriviaQA (Joshi et al., 2017), SearchQA (Dunn et al., 2017), HotpotQA (Yang et al., 2018), and NaturalQuestions (NQ (Kwiatkowski et al., 2019)).

- Commonsense reasoning on RAIN-BOW (Lourie et al., 2021) including α NLI (Bhagavatula et al., 2020), CosmosQA (Huang et al., 2019), HellaSWAG (Zellers et al., 2019), PIQA (Bisk et al., 2020), SocialIQA (Sap et al., 2019), and WinoGrande (Sakaguchi et al., 2020).
- Machine translation, including WMT EnDe (Bojar et al., 2014), WMT EnFr (Bojar et al., 2015), and WMT EnRo (Bojar et al., 2016).
- Summarization, including Aeslc (Zhang and Tetreault, 2019), BillSum (Kornilova and Eidelman, 2019), CNN/Dailymail (Hermann et al., 2015; See et al., 2017), Wikilingua (Ladhak et al., 2020), Gigaword (Graff et al., 2003; Rush et al., 2015), MultiNews (Fabbri et al., 2019), Newsroom (Grusky et al., 2018), SAMSum (Gliwa et al., 2019), and XSum (Narayan et al., 2018).
- Natural language generation on GEM (Gehrmann et al., 2021), including CommonGen (Lin et al., 2020), DART (Nan et al., 2021), E2E (Dušek et al., 2019), SGD (Rastogi et al., 2020), WebNLG (Gardent et al., 2017), WikiAuto (Jiang et al., 2020a), XSum, and Wikilingua.

C Task transferability results

The full results of our task transferability experiments can be found in Table 4. We show that in many cases, initializing the prompt to that of a source task can provide significant gain on a target task. Table 5 displays positive transfers with more than 10% relative error reduction on the target task.

D Task embedding similarity

In Figure 6, we show a clustered heatmap of cosine similarities between the task embeddings of the 26 NLP tasks we study in our task transferability experiments. For each task, we include the resulting task embeddings from all the three different prompt tuning runs on the task. As can be seen, our task embeddings capture task relationships: similar tasks

are grouped together into clusters. Additionally, task embeddings that are derived from different prompts of the same task are linked together.

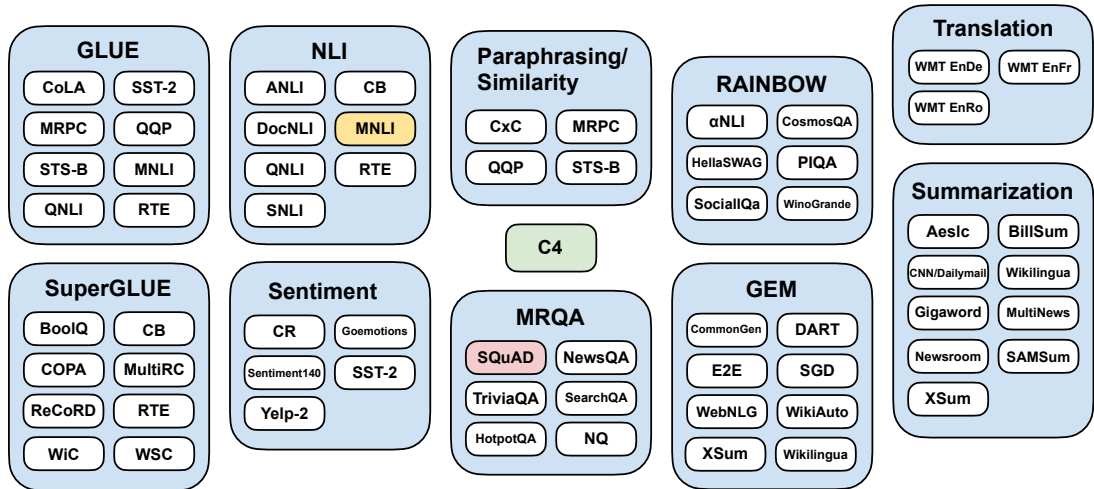


Figure 5: Datasets used in our PROMPTTUNING with SPOT experiments in Section 2. C4, MNLI, and SQuAD were all used by themselves as single source tasks in addition to being mixed in with other tasks.

	CoLA	STS-B	CR	MRPC	RTE	BoolQ	WiC	WSC	COPA	CB
Baseline	52.9 _{1.2}	88.1 _{0.6}	93.5 _{0.2}	86.1 _{0.7}	68.7 _{1.2}	73.0 _{1.2}	63.6 _{1.6}	71.5 _{1.7}	56.7 _{1.7}	92.7 _{1.9}
C4	54.8 _{1.1}	87.8 _{0.6}	93.9 _{0.1}	88.0 _{0.6}	69.1 _{1.9}	75.8 _{0.5}	66.3 _{0.8}	68.0 _{0.5}	54.3 _{0.9}	83.1 _{5.7}
DocNLI	52.7 _{0.9}	87.3 _{0.9}	93.6 _{0.4}	86.2 _{0.8}	67.4 _{2.6}	72.7 _{1.4}	64.7 _{0.3}	71.1 _{3.6}	56.0 _{5.9}	87.2 _{1.7}
Yelp-2	53.9 _{0.2}	88.1 _{0.3}	93.8 _{0.3}	86.6 _{0.8}	69.2 _{1.1}	74.8 _{0.7}	64.7 _{0.5}	70.8 _{1.2}	55.0 _{0.0}	87.8 _{1.6}
MNLI	54.2 _{0.7}	89.5 _{0.3}	93.9 _{0.4}	88.4 _{0.6}	74.7 _{1.3}	77.6 _{0.4}	69.5 _{0.5}	71.8 _{3.3}	69.3 _{2.1}	97.0 _{1.1}
QQP	55.6 _{1.3}	89.4 _{0.2}	93.7 _{0.5}	88.1 _{0.7}	72.0 _{0.5}	75.9 _{0.5}	67.9 _{0.2}	71.5 _{0.9}	62.0 _{2.2}	88.7 _{4.2}
QNLI	55.2 _{2.0}	89.2 _{0.2}	93.8 _{0.2}	87.8 _{0.1}	71.1 _{0.8}	75.6 _{0.5}	69.6 _{1.3}	71.5 _{2.5}	59.7 _{3.9}	92.5 _{1.1}
ReCoRD	54.7 _{1.3}	87.7 _{0.7}	93.7 _{0.1}	88.7 _{0.3}	67.5 _{1.3}	73.1 _{0.9}	65.5 _{0.9}	77.2 _{2.3}	59.3 _{1.2}	74.1 _{5.2}
CxC	55.0 _{0.2}	90.0 _{0.0}	93.9 _{0.2}	88.0 _{0.4}	70.3 _{0.5}	75.9 _{0.4}	70.2 _{0.1}	68.6 _{2.5}	60.3 _{3.9}	89.3 _{2.4}
SQuAD	54.9 _{1.2}	87.6 _{0.1}	93.9 _{0.5}	88.7 _{0.7}	71.2 _{0.4}	76.0 _{0.7}	66.8 _{0.3}	72.4 _{0.5}	63.0 _{1.6}	91.3 _{1.3}
DROP	53.0 _{1.0}	86.9 _{0.9}	93.7 _{0.2}	88.2 _{0.3}	65.7 _{3.1}	73.6 _{1.3}	67.5 _{1.2}	73.4 _{2.0}	60.0 _{3.6}	78.5 _{8.6}
SST-2	52.3 _{0.3}	87.9 _{0.3}	93.8 _{0.5}	85.6 _{0.9}	66.9 _{1.1}	73.3 _{0.5}	63.8 _{1.7}	68.6 _{0.4}	57.0 _{2.2}	92.9 _{1.3}
WinoGrande	52.8 _{1.6}	87.8 _{0.3}	93.7 _{0.1}	86.1 _{0.5}	67.9 _{1.3}	74.1 _{0.8}	62.4 _{2.5}	71.5 _{2.5}	56.7 _{1.2}	83.9 _{0.8}
HellaSWAG	32.7 _{23.6}	87.5 _{0.2}	93.6 _{0.0}	86.6 _{1.4}	63.9 _{5.4}	70.0 _{2.6}	60.1 _{3.9}	70.2 _{2.1}	58.0 _{2.2}	85.5 _{2.6}
MultiRC	50.0 _{4.6}	88.2 _{0.2}	93.4 _{0.1}	86.4 _{1.3}	67.6 _{1.0}	74.0 _{0.5}	66.4 _{0.5}	69.2 _{4.1}	56.0 _{4.1}	80.0 _{8.6}
CosmosQA	52.1 _{2.3}	87.7 _{0.5}	93.6 _{0.3}	87.9 _{0.8}	68.7 _{1.6}	73.4 _{1.3}	65.9 _{1.0}	69.6 _{3.2}	62.3 _{5.0}	83.9 _{8.8}
RACE	52.5 _{2.8}	87.5 _{0.5}	93.4 _{0.2}	86.5 _{0.8}	66.5 _{2.0}	73.6 _{0.5}	63.1 _{5.3}	68.9 _{1.2}	57.3 _{1.2}	84.8 _{3.4}

Table 4: Tasks can benefit each other via their prompts. The orange-colored row shows the results of prompt tuning T5 BASE on the target tasks without any prompt transfer. Each cell in the other rows represents the target task performance when transferring the prompt from the associated source task (row) to the associated target task (column). Positive transfers are shown in green and the best results are highlighted in bold (green). Numbers in the subscript indicate the standard deviation across 3 random seeds.

Transfer	Increase (relative)
MNLI → CB	58.9
MNLI → COPA	29.1
ReCoRD → WSC	20.0
MNLI → RTE	19.2
ReCoRD → MRPC	18.7
SQuAD → MRPC	18.7
CxC → WiC	18.1
MNLI → BoolQ	17.0
MNLI → MRPC	16.5
QNLI → WiC	16.5
MNLI → WiC	16.2
CxC → STS-B	16.0
DROP → MRPC	15.1
SQuAD → COPA	14.5
QQP → MRPC	14.4
CxC → MRPC	13.7
C4 → MRPC	13.7
CosmosQA → MRPC	12.9
CosmosQA → COPA	12.9
QQP → COPA	12.2
QNLI → MRPC	12.2
QQP → WiC	11.8
MNLI → STS-B	11.8
SQuAD → BoolQ	11.1
QQP → STS-B	10.9
QQP → BoolQ	10.7
CxC → BoolQ	10.7
DROP → WiC	10.7
QQP → RTE	10.5
C4 → BoolQ	10.4

Table 5: Positive transfers with more than 10% relative error reduction on the target task. $s \rightarrow t$ denotes the transfer from source task s to target task t .




Figure 6: Our prompt-based task embeddings capture task relationships: similar tasks are grouped together into clusters. Additionally, task embeddings that are derived from different prompts of the same task are linked together. t_1, t_2, t_3 correspond to three different prompt tuning runs on task t .