

# RIFF: Learning to Rephrase Inputs for Few-shot Fine-tuning of Language Models

Anonymous ACL submission

## Abstract

Pre-trained Language Models (PLMs) can be accurately fine-tuned for downstream text processing tasks. Recently, researchers have introduced several parameter-efficient fine-tuning methods that optimize input prompts or adjust a small number of model parameters (e.g LoRA). In this study, we explore the impact of altering the *input text* of the original task in conjunction with parameter-efficient fine-tuning methods. To most effectively rewrite the input text, we train a few-shot paraphrase model with a Maximum-Marginal Likelihood objective. Using six few-shot text classification datasets, we show that enriching data with paraphrases at train and test time enhances the performance beyond what can be achieved with parameter-efficient fine-tuning alone.

## 1 Introduction

Multiple Pre-trained Language Models (PLMs), such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), T5 (Raffel et al., 2019), and GPT2 (Radford et al., 2019b), have demonstrated remarkable performance when fine-tuned for downstream text processing tasks. PLM variants with less than 1 billion parameters are easier to train end-to-end with commodity hardware. However, very recent PLMs have been trained with a few hundred billion parameters, including PaLM-2 (540B)(Anil et al., 2023), GPT3 (175B)(Brown et al., 2020a), OPT (175B)(Zhang et al., 2022a), or Llama-2 (70B)(Touvron et al., 2023). Training all parameters of these models end-to-end is not straightforward unless done with a dedicated cluster with specialized hardware.

In response, NLP research have developed effective techniques to control or alter the behavior of PLMs by updating the input context through prompt optimization (Liu et al., 2021a) or adapting a few additional parameters within the network itself (Hu et al., 2021). However, current PLM

control techniques have not considered altering the *original input text* to improve the performance of the model. Here, we investigate this idea by training a secondary smaller PLM to paraphrase the original input at train and test time, thus augmenting the existing data and improving model performance.

Our inspiration comes from interactions with young children. Determining what a child knows is challenging because they can be sensitive to the wording of the question (Donaldson, 1978). Adults are also influenced by different wordings of a question. For example, opinion polling has been found to be sensitive to the wording of questions (Broughton, 1995). Just like we rephrase questions for humans, we should consider rephrasing input text while querying a PLM. For instance, while classifying the topic of a sentence, phrases related to time may be irrelevant and could be removed to simplify the modeling problem. Slight changes to wording could result in the model producing a correct prediction.

We explore the integration of paraphrased input texts during both the training and testing phases. At *training time*, augmenting data through paraphrase generation has been shown to enhance performance while updating all parameters of the model (Wei and Zou, 2019; Feng et al., 2021; Chen et al., 2021; Abaskohi et al., 2023). We broaden the scope of previous investigations by using paraphrase augmentation in tandem with recent prompt optimization and efficient tuning methods. At *test time*, recent works have used ensemble predictions with various optimized prompts and tuned weights (Izmailov et al., 2019; Li et al., 2023). We further contribute to this line of work by incorporating ensemble predictions based on input paraphrases, again in concert with prompt optimization and efficient tuning techniques.

We start by pre-training a smaller language model on paraphrases generated by a large lan-

guage model (i.e. ChatGPT or GPT3.5-turbo). Subsequently, we explore various training objectives for fine-tuning this paraphrase generator with feedback from the main task’s language model. Our analysis shows that our proposed objective reduces hallucination in the generated paraphrases. Then, we experiment on six text classification datasets demonstrating that incorporating paraphrase augmentation during both training and testing phases enhances the performance of discrete/soft prompt optimization and efficient tuning techniques. In summary, our contributions are as follows:

- We propose an efficient idea for Rephrasing Inputs for parameter-efficient Few-shot Fine-tuning of language models (RIFF) tested with recent prompt optimization and efficient tuning methods.
- We conduct a comprehensive study on various learning objectives for fine-tuning a paraphrase generator using feedback from the main language model.
- Our augmentation experiments on six text classification datasets reveal that paraphrase generation, when combined with prompt optimization and adaptation techniques is a simple yet effective approach to boost performance.

## 2 Problem Formulation

We focus on classification problems in Natural Language Understanding (NLU) tasks where we have access to a mini-batch of supervised training examples  $B_{\text{supp}} = \{(x_i, y_i)\}_{i=1}^N$ . Our goal is to update the parameter set  $\theta_{\text{lm}}$  for a language model by maximizing the probability of the class label  $y_i$  given the input  $x_i$ :  $P_{\theta_{\text{lm}}}(y_i|x_i)$ . Here, we augment  $B_{\text{supp}}$  with semi-supervised examples. In particular, we generate  $M$  paraphrases for each  $x_i$  using the paraphrase generator  $P_{\theta_{\text{par}}}(z_{i,j}|x_i)$ , where  $z_{i,j}$  represents the  $j$ -th paraphrase for the input  $x_i$ . In the optimal case, this paraphrase will preserve semantic meaning but vary syntactic/lexical form. We then incorporate the generated paraphrases to create a new mini-batch of examples  $B_{\text{s+p}} = B_{\text{supp}} \cup B_{\text{para}}$ . Using this augmented mini-batch, we optimize the following objective  $J_{\theta_{\text{lm}}}$ :

$$\sum_{i=1}^N \left\{ \log P_{\theta_{\text{lm}}}(y_i|x_i) + \frac{1}{M} \sum_{j=1}^M \log P_{\theta_{\text{lm}}}(y_i|z_{i,j}) \right\} \quad (1)$$

To train the language model using Equation 1, we need to update the parameter set  $\theta_{\text{lm}}$ . One ap-

proach would involve updating every parameter for the language model to optimize the training objective (referred to here as the "All-Finetuning" or *AllTune* approach). However, this method can be computationally intensive. As a result, we will explore the impact of paraphrase augmentation along with six other efficient baseline tuning techniques (Houlsby et al., 2019a) and prompt optimization (Liu et al., 2021b).

We assume that each input  $x$  or its paraphrase  $z$  is preceded by the task instruction  $p$ , which is often specified in previous works. The task instruction, which we represent using the symbol  $p$  to be consistent with prompt optimization literature, serves as a parameter-free, gradient-free technique for enhancing the performance of the PLM across various downstream tasks (Brown et al., 2020b; Petroni et al., 2019; Deng et al., 2022). When using only the task instructions, no parameters for the language model are updated ( $\theta_{\text{lm}} = \emptyset$ ), and zero-shot predictions are made solely on the evaluation data. We further investigate multiple language model tuning techniques while incorporating these task instructions into the input or its paraphrases.

### 2.1 LM-Friendly Paraphrase Search

Given a training example  $(x, y)$ , our objective is to assign the gold label  $y$  to the input  $x$  by maximizing the log likelihood  $\log P(y|x)$ . We leverage the fact that when  $x$  is misclassified, there may exist paraphrases of the input  $x$  that lead to the correct class prediction. These paraphrases should retain the semantic meaning of  $x$  while exhibiting syntactic differences, akin to the way we rephrase things when we have been misunderstood. Thus, we generate paraphrases  $z_j$  based on the input  $x$ , that enable the downstream language model to predict the correct label  $y$  with greater confidence. Consequently, our data log likelihood is factorized into the following marginalization over the space of paraphrases, where  $\theta_{\text{par}}$  and  $\theta_{\text{lm}}$  represent the parameters for the paraphrase generator and the downstream language model, respectively:

$$\begin{aligned} J_{\theta_{\text{par}}} &:= \log P(y|x) = \log E_z[P(y, z|x)] \\ &= \log \sum_z P_{\theta_{\text{par}}}(z|x) \times P_{\theta_{\text{lm}}}(y|z) \quad (2) \end{aligned}$$

To train the paraphrase generator and optimize the objective stated in Equation 2, we explore four distinct learning aspects: (a) two methods for gradient approximation, (b) a reward normalization

technique, (c) three decoding techniques for sampling paraphrases, and (d) two approaches to ensure grammatical integrity during paraphrase generation. By combining these elements, we examine various learning approaches to refine the paraphrase generator with the aid of the downstream language model. In the subsequent paragraphs, we will describe our suggested options for each aspect.

**Gradient Approximation:** Text generation can be reformulated as an episodic reinforcement learning problem where an agent (i.e. a paraphrase generator) generates tokens (i.e. takes actions) one step at a time until reaching the end of the episode (i.e. selecting the end of sequence token). Therefore, for a given training example  $(x, y)$  and its paraphrase  $z$ , we define the terminal reward (i.e. goodness) for  $z$  as  $R(z) = \log P_{\theta_{\text{lm}}}(y|z)$ . When approximating the gradient vector of objective 2 concerning  $\theta_{\text{par}}$ , we propose two strategies. These include: (i) *Maximum Marginal Likelihood (MML)* and (ii) approximating the gradient vector of the paraphrase model via the *Policy Gradient (PG)* theorem. Notably, gradient updates using these two methods exhibit a close relationship, with the main difference lying in the posterior coefficient utilized to score each sample (Guu et al., 2017). We can recast the main objective presented in equation 2 into the following equation representing the expected reward:

$$J_{\theta_{\text{par}}} := \log E_{z \sim P_{\theta_{\text{par}}}(\cdot|x)}[e^{R(z)}] \quad (3)$$

Given each input  $x$ , if we extract paraphrase samples from  $P_{\theta_{\text{par}}}(\cdot|x)$  and approximate the expectation in  $J_{\theta_{\text{par}}}$  via numerical summation, we optimize the objective using MML estimation. This process results in the following gradient update:

$$\begin{aligned} \nabla J_{\theta_{\text{par}}}^{\text{MML}} &:= \nabla_{\theta_{\text{par}}} \log E_z[e^{R(z)}] \\ &= \sum_{j=1}^M \phi^{\text{MML}}(z_j) \times \nabla_{\theta_{\text{par}}} \log P_{\theta_{\text{par}}}(z_j|x) \\ \phi^{\text{MML}}(z_j) &= \frac{P_{\theta_{\text{par}}}(z_j|x) \times e^{R(z_j)}}{\sum_{j'=1}^M P_{\theta_{\text{par}}}(z_{j'}|x) \times e^{R(z_{j'})}} \quad (4) \end{aligned}$$

By introducing the log inside the expectation (applying Jensen’s inequality), we can optimize a surrogate lower bound for the objective presented in equation 3, resulting in the following policy

gradient approximation (Sutton et al., 1999):

$$\begin{aligned} \nabla J_{\theta_{\text{par}}}^{\text{PG}} &:= \nabla_{\theta_{\text{par}}} E_z[R(z)] \\ &= \sum_{j=1}^M \phi^{\text{PG}}(z_j) \times \nabla_{\theta_{\text{par}}} \log P_{\theta_{\text{par}}}(z_j|x) \\ \phi^{\text{PG}}(z_j) &= P_{\theta_{\text{par}}}(z_j|x) \times R(z_j) \quad (5) \end{aligned}$$

**Reward Normalization:** For our secondary learning aspect, we can either utilize the basic reward, denoted as  $R(z_j)$ , or normalize the rewards among the paraphrases of a given input  $x$ . This process of normalization is particularly useful because it prevents the training of the paraphrase generator with rewards of varying magnitudes, as different training examples are not equally challenging for the language model. Prior research suggests that such normalization of rewards can significantly enhance the performance of text generators across a variety of tasks (Guo et al., 2022). The normalized reward  $R^n$  is defined as follows:

$$\begin{aligned} R^n(z_j) &= \frac{R(z_j) - \mu_j}{\sigma_j}, \mu_j = \frac{1}{M} \sum_{j=1}^M R(z_j) \\ \sigma_j^2 &= \frac{1}{M} \sum_{j=1}^M (R(z_j) - \mu_j)^2 \quad (6) \end{aligned}$$

**Decoding Techniques:** To train the paraphrase generator, we use both the MML and PG gradient estimations which necessitates drawing  $M$  samples from the paraphrase generator. We implement three decoding techniques for this purpose. Firstly, we utilize diverse beam search decoding (Vijayakumar et al., 2018) to gather these  $M$  paraphrases. Secondly, in order to thoroughly explore the paraphrase space, we alternatively collect the  $M$  paraphrases using nucleus (top-p) sampling (Holtzman et al., 2020). For the top-p sampling, we establish a sampling threshold of  $p = 0.99$ , at which we collect the minimal set of tokens from the vocabulary with a cumulative probability of at least 0.99. We then re-sample tokens from this set. And thirdly, during the training phase we blend diverse beam search and top-p sampling. Here, we initially sample  $M$  paraphrases using both methods, then combine the top  $M/2$  samples from each output to construct our final  $M$  samples. During the test phase, we only use diverse beam search.

**Grammatical Integrity:** We describe three distinct modeling techniques for both the MML and

PG gradient estimations: On-policy learning, Off-policy learning and Proximal Policy Optimization (PPO).

As we are sampling paraphrases from  $P_{\theta_{\text{par}}}(z_j|x)$  and updating  $\theta_{\text{par}}$  using these samples, the paraphrase generator may start generating ungrammatical text during this default on-policy learning setting. Similar instances of degenerate generation have been reported in tasks like question generation (Najafi and Fyshe, 2023) and program synthesis (Liang et al., 2018).

To mitigate degenerate paraphrase generation, we experiment with off-policy sampling. Here, we maintain a fixed sampling module  $P_{\text{fixed}}(z_j|x)$  for sample selection, then update the main paraphrase generator  $P_{\theta_{\text{par}}}(z_j|x)$  within the frameworks of objectives 4 and 5. Consequently, with these off-policy samples, the posterior coefficients incorporate the importance sampling ratio  $s(z_j) = \frac{P_{\theta_{\text{par}}}(z_j|x)}{P_{\text{fixed}}(z_j|x)}$

$$\begin{aligned} \phi_{\text{off}}^{\text{PG}}(z_j) &= s(z_j) \times R(z_j) \\ \phi_{\text{off}}^{\text{MML}}(z_j) &= \frac{s(z_j) \times e^{R(z_j)}}{\sum_{j'=1}^M s(z_{j'}) \times e^{R(z_{j'})}} \quad (7) \end{aligned}$$

Our next solution for degenerate paraphrases involves imposing a penalty in the training objective if the samples drawn from the current paraphrase generator,  $P_{\theta_{\text{par}}}(z|x)$ , deviate from those of the pre-trained paraphrase generator. We can implement this penalty as a KL-divergence penalty between the distributions of paraphrases produced by the current model and the pre-trained one. This approach resembles the PPO learning with a KL penalty (Schulman et al., 2017) and has been used in fine-tuning InstructGPT with a reward model trained over human feedback (Ouyang et al., 2022). In the case of InstructGPT, researchers prevent the reward fine-tuned model from diverging from a separate language model pre-trained on supervised data (Ouyang et al., 2022). To integrate this penalty we define the following new objective for  $\theta_{\text{par}}$ :

$$\begin{aligned} J_{\theta_{\text{par}}}^{\text{PPO}} &:= \log E_z[e^{R(z)}] - \beta E_z[\log s(z)] \\ s(z) &= \frac{P_{\theta_{\text{par}}}(z|x)}{P_{\text{fixed}}(z|x)} \quad (8) \end{aligned}$$

Building upon the previously approximated MML and PG gradients, we can now derive the following regularized gradient vector with respect

to  $\theta_{\text{par}}$ . Please note that  $\beta$  is a hyper-parameter in this context:

$$\begin{aligned} \nabla J_{\theta_{\text{par}}} - \beta E_z[(\log s(z) + 1)\nabla \log P_{\theta_{\text{par}}}(z|x)] \\ z \sim P_{\theta_{\text{par}}}(\cdot|x) \quad (9) \end{aligned}$$

Note that the KL penalty can be interpreted as the sum of a grammar reward, denoted by  $\log P_{\text{fixed}}(z|x)$ , and an entropy regularization term over  $P_{\theta_{\text{par}}}(z|x)$ . The entropy regularization aids in the diverse exploration of the search space (Mnih et al., 2016), while the grammar reward discourages the learning of ungrammatical samples.

## 2.2 Ensemble Inference

After optimizing Equation 2 and fine-tuning our paraphrase generator, we generate weakly-supervised examples for inclusion in Equation 1 to train our downstream language model.

To predict the class label of a test example, we could either use our fine-tuned language model to predict the class label based on the original input  $x$ , or adopt an ensemble approach. For the latter, for a given  $x$ , we generate  $M$  paraphrases using our fine-tuned paraphrase generator. We then average the prediction scores for a potential class label across the  $M+1$  values according to Equation 1 to predict the class label for that input example  $x$ . This aligns with our earlier assumption that some paraphrases could be easier for the language model to predict the correct class label. During data augmentation for the language model, we select the validation set’s best model according to this ensemble prediction.

## 3 Experiments

### 3.1 Setup

#### Pre-trained Models:

For paraphrase generation, we employ a T5-base model (Raffel et al., 2019) which has been trained on paraphrases generated by ChatGPT (i.e. version GPT3.5-turbo). These output paraphrases were generated for input texts from various datasets, including Quora paraphrase questions, texts from SQUAD 2.0, and the CNN news dataset (Vladimir Vorobev, 2023). To create this training data, ChatGPT generated five paraphrases for each input, which were then used as the target for the T5-base model. The weights for this model are publicly available<sup>1</sup>. In our experiments,

<sup>1</sup>[https://huggingface.co/humarin/chatgpt\\_paraphraser\\_on\\_T5\\_base](https://huggingface.co/humarin/chatgpt_paraphraser_on_T5_base)

this model was able to generate more diverse paraphrases compared to other public pre-trained models.

For our main language model, we use the RoBERTa-large model pre-trained with the Masked Language Modeling (MLM) objective (Liu et al., 2019), which has demonstrated strong performance on NLU tasks. Our proposed learning framework can be readily extended to other paraphrase generators or backbone language models.

**Datasets:** Inspired by prior work (Gao et al., 2021; Deng et al., 2022), we experiment on six classification tasks in the few-shot setting. These include sentiment classification tasks such as the binary sentiment datasets SST2 (Socher et al., 2013), CR (Hu and Liu, 2004), and MR (Pang and Lee, 2005). We also experiment on the 5-label sentiment dataset SST5 (Socher et al., 2013), the question type classification dataset TREC (Voorhees and Tice, 2000), and the topic classification dataset AG-News (Zhang et al., 2015). The number of classes per dataset, as well as the used instructions are outlined in Appendix G. Instructions and class verbalizers are based on previous work (Deng et al., 2022) in prompt optimization. Detailed information about the specific learning rates for each LM technique along with other hyper-parameters can be found in Appendix D.

### 3.2 Few-shot Paraphrase Fine-Tuning

As discussed in Section 2.1, there are four learning aspects to be considered when fine-tuning our paraphrase generator for the downstream language model. We conduct an extensive set of experiments in the 128-shot setting for the SST2 binary sentiment classification task.




Figure 1: Average ensemble accuracy over five validation splits in the 128-shot SST2 classification task. PG gradient estimation is not robust during the training trajectory while doing on-policy learning.

We randomly select 128 training examples for each unique label within the dataset. An equal number of examples are gathered to form an internal validation set. We create five train/validation splits using the arbitrarily chosen random seeds. We train the models for 1120 training steps with the batch size of 8 (i.e. 35 epochs). As we are training the models, we evaluate the performance of 140 weight checkpoints per model on the validation splits (i.e. one checkpoint per 8 training steps). We examine the mean accuracy, which is averaged over the five validation splits. Despite the ensembling approach described in Section 2.2, to accurately capture the quality of the generated paraphrases, we exclude the original input  $x$  when computing the ensemble accuracy on the validation splits.

We assess the impact of reward normalization in the context of on-policy, off-policy, and PPO learning, considering both PG and MML gradient estimations. Table 1 lists the best performance out of all the checkpoints evaluated on the validation splits, which is further averaged over five validation splits. With both PG and MML gradient estimations, reward normalization is boosting the performance across the three text decoding techniques for both on-policy and PPO learning techniques (see ‘AVG’ column in Table 1). Conversely, reward normalization is not improving performance with off-policy learning (follow discussion in Appendix B and see Table 4)

Table 1 verifies that MML gradient estimation outperforms PG gradient estimation on average across three decoding techniques for both on-policy and PPO learning techniques. The highest accuracy is achieved by ‘PG-Z’ with on-policy learning and top-p decoding, however it is not robust during the entire training trajectory. Figure 1 shows that PG gradient estimation is not robust throughout the training trajectory, which causes the paraphrase generator to produce nonsensical paraphrases. This results in downstream classification performance on par with random guessing. In contrast, off-policy and PPO learning circumvent this divergence. MML gradient estimation maintains robustness throughout the training phase. In Table 1, we also report the average accuracy of all the checkpoints as we are training the models (numbers in parentheses). The learning technique ‘MML-Z’ is more robust during the training trajectory compared to ‘PG-Z’.

Upon investigating various elements of our learning objectives for fine-tuning the paraphrase gener-

Table 1: The average accuracy of the best performing validation checkpoint in the 128-shot SST2 classification task for both the on-policy and PPO learning techniques. Highest performance per column bolded. Last column reports the macro-average among each row. Numbers in parentheses ( $\sigma$  |  $m$ ):  $\sigma$  represents the standard deviations for the reported means across five train/validation splits;  $m$  reports the average accuracy of all the validation checkpoints as we monitor robustness during the training trajectory. The suffix ‘-Z’ denotes models trained with reward normalization.

Learning Technique	On-Policy			PPO			AVG
	Top-P	Beam	Mixed	Top-P	Beam	Mixed	
No Tuning	67.5	67.5	67.5	67.5	67.5	67.5	67.5
PG	67.9 (2.3   53.2)	68.0 (1.4   52.0)	67.9 (2.0   52.4)	68.5 (1.2   67.4)	68.3 (1.9   67.4)	69.1 (1.4   68.2)	68.3 (1.7   60.1)
PG-Z	<b>71.3</b> (2.1   63.8)	<b>70.2</b> (1.7   68.6)	<b>71.2</b> (1.3   66.6)	68.9 (1.3   67.6)	68.8 (1.4   67.9)	69.8 (1.3   68.5)	70.0 (1.5   67.2)
MML	69.6 (2.1   68.5)	69.1 (1.8   67.6)	69.8 (2.0   68.6)	<b>69.5</b> (2.4   68.2)	69.9 (1.9   69.0)	70.5 (3.0   68.9)	69.7 (2.2   68.5)
MML-Z	70.3 (2.7   69.1)	<b>70.2</b> (2.0   68.9)	70.2 (2.3   69.1)	68.9 (1.8   67.9)	<b>70.3</b> (1.7   69.0)	<b>70.6</b> (2.5   68.9)	<b>70.1</b> (2.2   68.8)

ator, the combination that delivers the best performance across the validation splits, which is also robust during the entire training trajectory, includes: **MML** gradient approximation, **PPO** learning, **mixed decoding** for sample generation, and finally applying **reward normalization**. We name this combined approach our proposed RIFF algorithm.

### 3.3 Paraphrase Quality Analysis

We investigate the impact of the RIFF algorithm on the quality of the paraphrases in the 128-shot setting across three classification tasks SST2, SST5 and AGNews. To evaluate the quality of the paraphrases, we report the following five metrics:

**Grammar (GR):** We evaluate grammar by calculating the perplexity score averaged across the dataset using the GPT-2-Large model (Radford et al., 2019a). A low *GR* score indicates more grammatical texts.

**Lexical Diversity (LD):** To assess how paraphrases differ lexically from the *original* input text, we calculate the unigram and bigram Rouge scores between each paraphrase and the original input text (Lin, 2004). We then report  $1 - \frac{(Rouge_1 + Rouge_2)}{2}$  as our lexical diversity metric. A higher *LD* score indicates greater lexical difference compared to the original text.

**Pair-wise Lexical Diversity (PLD):** To assess the lexical diversity among the set of *paraphrases* for a given original input text, we calculate *LD* scores for every pair of paraphrases for an input text, and report the average. A higher *PLD* score indicates greater diversity among the paraphrases for a specific input text.

**Semantic Similarity (SE):** To assess how semantically similar paraphrases are to the original in-

Table 2: Average metrics on the test sets to evaluate the quality of paraphrases across three classification tasks: SST2, SST5, and AGNews datasets. The metrics are further averaged across five training folds for the *FinPara* method. *OrigIn*: Represents the original task inputs. *PrePara*: Corresponds to task inputs obtained from the pre-trained paraphraser. *FinPara*: Indicates task inputs from the finetuned paraphraser in the 128-shot setting. We scale scores into the range of [0-100], except for *GR*. Better performance per column bolded

Input Type	<i>GR</i>	<i>LD</i>	<i>PLD</i>	<i>SE</i>	<i>FC</i>
<i>OrigIn</i>	198	n/a	n/a	n/a	n/a
<i>PrePara</i>	<b>143</b>	<b>60.5</b>	<b>61.6</b>	70.4	77.3
<i>FinPara</i>	162	50.6	53.3	<b>74.7</b>	<b>79.1</b>

put text, we compute the BERTScore’s F1 metric (Zhang\* et al., 2020) between each paraphrase and the original input text using the BERT-Large model. A higher *SE* score signifies more semantic similarity with the original text.

**Factual Consistency (FC):** To measure hallucination in the generated paraphrases with respect to the original input text, we rely on a publicly available factual consistency metric. The model has been trained for textual entailment and summarization datasets with samples annotated for factual consistency<sup>2</sup>.

We present the metrics for the datasets in Table 2. Compared to a model that was not fine-tuned for this task (PrePara), our RIFF algorithm has reduced the diversity among the generated paraphrases and their lexical variation compared to the original input text. This outcome aligns with our search-learn objective that prioritizes high-scoring paraphrases over others. RIFF has contributed

<sup>2</sup>[https://huggingface.co/vectara/hallucination\\_evaluation\\_model](https://huggingface.co/vectara/hallucination_evaluation_model)

higher semantic similarity compared to the original input. Interestingly, the perplexity of the generated paraphrases after fine-tuning with our objective is still low, demonstrating the grammatical accuracy of these paraphrases. The example paraphrases shown in Appendix C illustrate that RIFF reduces hallucination in the generated paraphrases, which may contribute to the lower *LD* score but higher *SE* score with respect to the original input text. A higher *FC* score as shown in Table 2 verifies that the RIFF objective has reduced hallucination in the generated paraphrases.

### 3.4 Paraphrases for Few-shot LM Tuning

Our primary hypothesis is that various LM tuning techniques could benefit from diverse views of the original input text. To test this hypothesis, we fine-tuned our paraphrase generators in a 16-shot classification setup using the RIFF algorithm. Subsequently, we fine-tuned the downstream classification model in the same 16-shot setting, while introducing  $M = 8$  paraphrases as per the objective outlined in Equation 1. For evaluation, we used the best model from the validation set to make predictions on standard evaluation splits, following the ensemble approach described in Section 2.2. For consistency with prior research, we used the random dataset splits provided by RLPrompt (Deng et al., 2022), aligning with the random seeds used by LM-BFF (Gao et al., 2021).

We study the effect of paraphrases on seven language model tuning techniques. The first technique *AllTune* updates every parameter in the network. Another technique, *GS*, is based on AutoPrompt (Shin et al., 2020) for discrete prompt optimization. The technique *SpTune* (Lester et al., 2021) learns soft prompt vectors, and *LoRA* (Hu et al., 2021) is a recent adaptation technique. Additionally, we investigate *ClsTune*, which trains a classifier on top of the language model, *InTune*, which updates all of the input embedding table, and *HTune*, which only updates the language modeling head in the Transformer architecture. A detailed description of these techniques is discussed in Appendix A.

Table 3 illustrates the average accuracy on standard test sets across six text classification datasets. The reported scores correspond to four distinct LM tuning techniques: *GS*, *SpTune*, *AllTune*, and *LoRA*. Results for *ClsTune*, *HTune*, and *InTune* are presented in Table 11 of Appendix E.

Recent prompt optimization techniques like *GS*

and *SpTune* significantly benefit from paraphrase augmentation during training, with *SpTune* demonstrating the most dramatic improvement (2.2% average accuracy increase). While *LoRA* already outperforms these techniques (Hu et al., 2021), paraphrase augmentation further enhances its efficiency in learning adaptation matrices, leading to average accuracy gains of 0.2% on SST2 and 0.3% on AGNews. When coupled with ensemble predictions, denoted in rows with “+RIFF (train+test)”, all LM tuning techniques see improvements from the generated paraphrases.

### 3.5 Paraphrase Robustness Analysis




Figure 2: Average test accuracy across six classification datasets. Input at rank 0 represents the original test input, while the remaining eight inputs are top-ranked paraphrases generated by our fine-tuned paraphrase model.

We generate  $M = 8$  paraphrases for each test input text and investigate the average performance of the fine-tuned language models on each of these paraphrases compared to the original input text. Figure 2 illustrates the average test accuracy over six classification datasets for *AllTune*, *LoRA*, and *SPTune*, which are popular LM tuning techniques in NLP. The original input texts are denoted with inputs at rank 0, whereas rank  $i$  ( $i \in \{1, 2, 3, 4, 5, 6, 7, 8\}$ ) are the paraphrases returned by the diverse beam search. We observe that the LM tuning techniques are not robust to paraphrases. Our paraphrase augmentation during training increases paraphrase robustness. *SPTune* observes more significant improvement in robustness from paraphrase augmentation during training.

## 4 Related Works

To improve prompt optimization and efficient tuning techniques for LMs, we incorporate the generated paraphrases into the training mini-batches.

Table 3: Average accuracy on the standard evaluation sets for the 16-shot text classification. Numbers in parentheses are standard deviations across the five train/validation folds. The last column is the micro averaged performance across the datasets. Highest performance per dataset bolded, second highest underlined. †: the average 16-shot *AllTune* results with automatically searched templates (Gao et al., 2021). \*: reported zero-shot results for GPT3 with in-context learning (Gao et al., 2021).

Tuning Method	SST2	SST5	CR	MR	TREC	AGN	AVG
No Tuning	84.6	31.0	77.8	81.3	27.6	51.5	58.6
GPT-3*	84.8	30.6	87.4	80.5	26.2	-	66.9
RLPrompt	92.5	41.4	89.5	<u>87.1</u>	60.5	80.2	77.7
LM-BFF†	92.3	49.2	89.0	85.5	<u>88.2</u>	-	78.5
<i>GS</i>	85.5 (1.3)	37.0 (4.2)	80.2 (2.3)	83.0 (1.8)	45.3 (13.1)	82.0 (1.4)	75.0 (2.3)
+RIFF (train)	86.4 (1.9)	37.8 (3.3)	82.7 (1.3)	84.7 (2.1)	51.0 (8.6)	81.0 (2.4)	75.4 (2.5)
+RIFF (train+test)	87.3 (2.0)	38.2 (3.5)	85.1 (1.9)	84.7 (1.9)	52.4 (7.7)	83.3 (1.5)	77.0 (2.1)
<i>SpTune</i>	89.7 (3.7)	39.4 (6.2)	82.4 (2.8)	86.1 (2.2)	35.2 (2.7)	82.0 (2.6)	76.1 (3.2)
+RIFF (train)	91.2 (2.2)	44.5 (4.2)	84.6 (1.9)	86.1 (0.7)	38.4 (4.3)	84.0 (1.9)	78.3 (2.2)
+RIFF (train+test)	91.6 (2.3)	45.1 (4.1)	86.2 (1.8)	86.6 (0.8)	38.4 (4.0)	86.0 (1.0)	79.7 (1.7)
<i>AllTune</i>	93.1 (0.4)	48.0 (1.0)	89.2 (0.8)	<b>87.3</b> (3.1)	87.2 (3.8)	87.7 (0.5)	<u>83.0</u> (1.0)
+RIFF (train)	<u>93.6</u> (1.3)	<u>50.6</u> (1.0)	<u>90.2</u> (1.5)	85.8 (2.3)	84.2 (4.9)	87.2 (0.6)	<u>83.0</u> (1.2)
+RIFF (train+test)	<b>93.8</b> (1.2)	<b>51.2</b> (1.6)	<b>91.0</b> (1.6)	85.5 (2.3)	84.4 (4.9)	87.2 (0.6)	<b>83.2</b> (1.3)
<i>LoRA</i>	92.5 (1.8)	48.1 (1.8)	88.6 (2.0)	86.0 (2.6)	<b>89.3</b> (2.2)	87.3 (0.5)	82.6 (1.3)
+RIFF (train)	92.7 (1.8)	48.0 (2.3)	87.5 (1.5)	85.1 (2.9)	84.8 (2.7)	<u>87.6</u> (0.3)	82.3 (1.3)
+RIFF (train+test)	93.1 (1.2)	49.2 (2.0)	89.0 (1.1)	85.4 (2.8)	85.9 (3.6)	<b>87.9</b> (0.3)	82.9 (1.1)

Paraphrase generation represents just one technique of data augmentation. For a comprehensive overview of diverse data augmentation techniques for NLP tasks, we direct interested readers to a recent survey by Chen et al. (2021).

A recent work for few-shot prompt-based learning helps contrastive training by paraphrasing the inputs (Abaskohi et al., 2023). Our work proposes an objective to further fine-tune the paraphrase generator distilled from an LLM that reduces hallucination. Despite the previous work, which only studies the *AllTune* technique, we investigate the impact of paraphrases for various language model tuning techniques. Without contrastive learning, we can show that we can improve LM’s robustness by paraphrase generation during training.

**Prompt Optimization & Efficient Tuning:** Recent research proposes various techniques for prompt optimization and efficient tuning. We have used successful techniques from each of these areas. Appendix F provides our detailed description of these recent techniques. All of the recent techniques for prompt optimization and efficient tuning use the original input context provided within the dataset.

**Paraphrase Generation:** Recent techniques encompass various approaches, including the use of copy mechanisms, Variational Autoencoders,

Generative Adversarial Networks, and Reinforcement Learning techniques to generate diverse paraphrases (Zhou and Bhat, 2021). While previous studies have applied RL techniques for paraphrase generation, we propose the use of MML gradients instead of policy gradients to fine-tune our paraphrase model by the reward of a secondary classification task (see Appendix F for more discussion).

## 5 Conclusion

We investigated the impact of incorporating input paraphrases while fine-tuning PLMs with recent efficient tuning techniques. We also provided extensive experiments for reducing noise in a distantly supervised paraphrase generator.

## Limitations

Our paraphrase generator is pre-trained on semi-supervised paraphrases given by a truly large language model (i.e. ChatGPT). Although these large models are capable of generating high quality paraphrases for the English language. It is not clear if these semi-supervised paraphrases are available for other languages.

We hypothesize that our proposed objective can be used to generate paraphrases that minimize the downstream language model’s perfor-



638	mance on unseen data by maximizing the objective	
639	$\log(1 - P(y x))$ . Exploring the connection to Ad-	
640	versarial Generative Networks and textual attacks	
641	on language models through paraphrase generation	
642	presents a promising avenue for future work.	
643	<b>Ethics Statement</b>	
644	Many language models show biases in their output	
645	due to the data used to train them (Liang et al.,	
646	2021). It is possible that even with few-shot lan-	
647	guage model tuning, we might continue to detect	
648	analogous biases in the downstream classification	
649	task, for instance, resulting in diminished classi-	
650	fication accuracy for specific minority groups. It	
651	is also possible that the additional data generated	
652	by the paraphrase model will exaggerate existing	
653	biases.	
654	<b>References</b>	
655	Amirhossein Abaskohi, Sascha Rothe, and Yadollah	
656	Yaghoobzadeh. 2023. <a href="#">LM-CPPF: Paraphrasing-</a>	
657	<a href="#">guided data augmentation for contrastive prompt-</a>	
658	<a href="#">based few-shot fine-tuning</a> . In <i>Proceedings of the</i>	
659	<i>61st Annual Meeting of the Association for Computa-</i>	
660	<i>tional Linguistics (Volume 2: Short Papers)</i> , pages	
661	670–681, Toronto, Canada. Association for Computa-	
662	tional Linguistics.	
663	Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin John-	
664	son, Dmitry Lepikhin, Alexandre Passos, Siamak	
665	Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng	
666	Chen, Eric Chu, Jonathan H. Clark, Laurent El	
667	Shafey, Yanping Huang, Kathy Meier-Hellstern, Gau-	
668	rav Mishra, Erica Moreira, Mark Omernick, Kevin	
669	Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao,	
670	Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez	
671	Abrego, Junwhan Ahn, Jacob Austin, Paul Barham,	
672	Jan Botha, James Bradbury, Siddhartha Brahma,	
673	Kevin Brooks, Michele Catasta, Yong Cheng, Colin	
674	Cherry, Christopher A. Choquette-Choo, Aakanksha	
675	Chowdhery, Clément Crepy, Shachi Dave, Mostafa	
676	Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz,	
677	Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu	
678	Feng, Vlad Fienber, Markus Freitag, Xavier Gar-	
679	cia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-	
680	Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua	
681	Howland, Andrea Hu, Jeffrey Hui, Jeremy Hur-	
682	witz, Michael Isard, Abe Ittycheriah, Matthew Jagiel-	
683	ski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun,	
684	Sneha Kudugunta, Chang Lan, Katherine Lee, Ben-	
685	jamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li,	
686	Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu,	
687	Frederick Liu, Marcello Maggioni, Aroma Mahendru,	
688	Joshua Maynez, Vedant Misra, Maysam Moussalem,	
689	Zachary Nado, John Nham, Eric Ni, Andrew Nys-	
690	trom, Alicia Parrish, Marie Pellat, Martin Polacek,	
691	Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif,	
	Bryan Richter, Parker Riley, Alex Castro Ros, Au-	692
	rko Roy, Brennan Saeta, Rajkumar Samuel, Renee	693
	Shelby, Ambrose Slone, Daniel Smilkov, David R.	694
	So, Daniel Sohn, Simon Tokumine, Dasha Valter,	695
	Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang,	696
	Pidong Wang, Zirui Wang, Tao Wang, John Wiet-	697
	ing, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting	698
	Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven	699
	Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav	700
	Petrov, and Yonghui Wu. 2023. <a href="#">Palm 2 technical</a>	701
	<a href="#">report</a> .	702
	David Broughton. 1995. <i>The assumptions and theory</i>	703
	<i>of public opinion polling</i> , pages 15–33. Macmillan	704
	Education UK, London.	705
	Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie	706
	Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind	707
	Neelakantan, Pranav Shyam, Girish Sastry, Amanda	708
	Askell, Sandhini Agarwal, Ariel Herbert-Voss,	709
	Gretchen Krueger, Tom Henighan, Rewon Child,	710
	Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,	711
	Clemens Winter, Christopher Hesse, Mark Chen, Eric	712
	Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess,	713
	Jack Clark, Christopher Berner, Sam McCandlish,	714
	Alec Radford, Ilya Sutskever, and Dario Amodei.	715
	2020a. <a href="#">Language models are few-shot learners</a> .	716
	Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie	717
	Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind	718
	Neelakantan, Pranav Shyam, Girish Sastry, Amanda	719
	Askell, Sandhini Agarwal, Ariel Herbert-Voss,	720
	Gretchen Krueger, Tom Henighan, Rewon Child,	721
	Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,	722
	Clemens Winter, Christopher Hesse, Mark Chen,	723
	Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin	724
	Chess, Jack Clark, Christopher Berner, Sam Mc-	725
	Candlish, Alec Radford, Ilya Sutskever, and Dario	726
	Amodei. 2020b. <a href="#">Language models are few-shot learn-</a>	727
	<a href="#">ers</a> . <i>CoRR</i> , abs/2005.14165.	728
	Jiaao Chen, Derek Tam, Colin Raffel, Mohit Bansal, and	729
	Diyi Yang. 2021. <a href="#">An empirical survey of data aug-</a>	730
	<a href="#">mentation for limited data learning in NLP</a> . <i>CoRR</i> ,	731
	abs/2106.07499.	732
	Ganqu Cui, Wentao Li, Ning Ding, Longtao Huang,	733
	Zhiyuan Liu, and Maosong Sun. 2023. <a href="#">Decoder tun-</a>	734
	<a href="#">ing: Efficient language understanding as decoding</a> .	735
	In <i>Proceedings of the 61st Annual Meeting of the</i>	736
	<i>Association for Computational Linguistics (Volume 1:</i>	737
	<i>Long Papers)</i> , pages 15072–15087, Toronto, Canada.	738
	Association for Computational Linguistics.	739
	Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yi-	740
	han Wang, Han Guo, Tianmin Shu, Meng Song, Eric	741
	Xing, and Zhiting Hu. 2022. <a href="#">RLPrompt: Optimizing</a>	742
	<a href="#">discrete text prompts with reinforcement learning</a> .	743
	In <i>Proceedings of the 2022 Conference on Empiri-</i>	744
	<i>cal Methods in Natural Language Processing</i> , pages	745
	3369–3391, Abu Dhabi, United Arab Emirates. As-	746
	sociation for Computational Linguistics.	747
	Jacob Devlin, Ming-Wei Chang, Kenton Lee, and	748
	Kristina Toutanova. 2019. <a href="#">BERT: Pre-training of</a>	749

750	deep bidirectional transformers for language understanding.	Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski,	807
751	In <i>Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)</i> , pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.	Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019a. Parameter-efficient transfer learning for NLP.	808
752		In <i>Proceedings of the 36th International Conference on Machine Learning</i> , volume 97 of <i>Proceedings of Machine Learning Research</i> , pages 2790–2799. PMLR.	809
753			810
754			811
755			812
756			813
757	Margaret C. Donaldson. 1978. <i>Children's Minds</i> .		814
758	Wanyu Du and Yangfeng Ji. 2019. An empirical comparison on imitation learning and reinforcement learning for paraphrase generation.	Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski,	815
759	In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 6012–6018, Hong Kong, China. Association for Computational Linguistics.	Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019b. Parameter-efficient transfer learning for NLP. <i>CoRR</i> , abs/1902.00751.	816
760			817
761			818
762			819
763		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. <i>CoRR</i> , abs/2106.09685.	820
764			821
765			822
766			823
767	Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edward Hovy. 2021. A survey of data augmentation approaches for NLP.	Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews.	824
768	In <i>Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021</i> , pages 968–988, Online. Association for Computational Linguistics.	In <i>Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04</i> , page 168–177, New York, NY, USA. Association for Computing Machinery.	825
769			826
770			827
771			828
772			829
773	Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners.	Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. 2019. Averaging weights leads to wider optima and better generalization.	830
774	In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 3816–3830, Online. Association for Computational Linguistics.		831
775			832
776			833
777			834
778		Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners.	835
779		In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 22199–22213. Curran Associates, Inc.	836
780			837
781	Sonal Garg, Sumanth Prabhu, Hemant Misra, and G. Srinivasaraghavan. 2021. Unsupervised contextual paraphrase generation using lexical control and reinforcement learning.	Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning.	838
782	<i>CoRR</i> , abs/2103.12777.	In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	839
783			840
784			841
785	Han Guo, Bowen Tan, Zhengzhong Liu, Eric Xing, and Zhiting Hu. 2022. Efficient (soft) Q-learning for text generation with limited good data.		842
786	In <i>Findings of the Association for Computational Linguistics: EMNLP 2022</i> , pages 6969–6991, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.		843
787			844
788			845
789		Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation.	846
790		In <i>Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 4582–4597, Online. Association for Computational Linguistics.	847
791	Han Guo, Bowen Tan, Zhengzhong Liu, Eric P. Xing, and Zhiting Hu. 2021. Efficient (soft) q-learning for text generation with limited good data.		848
792			849
793			850
794			851
795			852
796			853
797		Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2023. Making language models better reasoners with step-aware verifier.	854
798		In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 5315–5333, Toronto, Canada. Association for Computational Linguistics.	855
799			856
800			857
801			858
802	Dan Hendrycks and Kevin Gimpel. 2016. Bridging nonlinearities and stochastic regularizers with gaussian error linear units.		859
803	<i>CoRR</i> , abs/1606.08415.		860
804		Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. Paraphrase generation with deep reinforcement learning.	861
805	Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The curious case of neural text degeneration.	In <i>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing</i> ,	862
806			863
			864

865	pages 3865–3878, Brussels, Belgium. Association for Computational Linguistics.	
866		
867	Chen Liang, Mohammad Norouzi, Jonathan Berant, Quoc V Le, and Ni Lao. 2018. <a href="#">Memory augmented policy optimization for program synthesis and semantic parsing</a> . In <i>Advances in Neural Information Processing Systems</i> , volume 31. Curran Associates, Inc.	
868		
869		
870		
871		
872		
873	Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2021. <a href="#">Towards Understanding and Mitigating Social Biases in Language Models</a> . In <i>International Conference on Machine Learning</i> , pages 6565–6576. PMLR.	
874		
875		
876		
877		
878	Chin-Yew Lin. 2004. <a href="#">ROUGE: A package for automatic evaluation of summaries</a> . In <i>Text Summarization Branches Out</i> , pages 74–81, Barcelona, Spain. Association for Computational Linguistics.	
879		
880		
881		
882	Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2020. <a href="#">Exploring versatile generative language model via parameter-efficient transfer learning</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 441–459, Online. Association for Computational Linguistics.	
883		
884		
885		
886		
887		
888	Mingtong Liu, Erguang Yang, Deyi Xiong, Yujie Zhang, Yao Meng, Changjian Hu, Jinan Xu, and Yufeng Chen. 2020. <a href="#">A learning-exploring method to generate diverse paraphrases with multi-objective deep reinforcement learning</a> . In <i>Proceedings of the 28th International Conference on Computational Linguistics</i> , pages 2310–2321, Barcelona, Spain (Online). International Committee on Computational Linguistics.	
889		
890		
891		
892		
893		
894		
895		
896		
897	Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. <a href="#">Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing</a> . <i>CoRR</i> , abs/2107.13586.	
898		
899		
900		
901		
902	Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021b. <a href="#">Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing</a> .	
903		
904		
905		
906	Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. <a href="#">P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks</a> . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)</i> , pages 61–68, Dublin, Ireland. Association for Computational Linguistics.	
907		
908		
909		
910		
911		
912		
913		
914	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. <a href="#">Roberta: A robustly optimized BERT pretraining approach</a> . <i>CoRR</i> , abs/1907.11692.	
915		
916		
917		
918		
919	Ilya Loshchilov and Frank Hutter. 2017. <a href="#">Fixing weight decay regularization in adam</a> . <i>CoRR</i> , abs/1711.05101.	
920		
921		
	Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. <a href="#">Asynchronous methods for deep reinforcement learning</a> . <i>CoRR</i> , abs/1602.01783.	922
		923
		924
		925
		926
	Saeed Najafi and Alona Fyshe. 2023. <a href="#">Weakly-supervised questions for zero-shot relation extraction</a> . In <i>Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics</i> , pages 3075–3087, Dubrovnik, Croatia. Association for Computational Linguistics.	927
		928
		929
		930
		931
		932
	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. <a href="#">Training language models to follow instructions with human feedback</a> .	933
		934
		935
		936
		937
		938
		939
		940
	Bo Pang and Lillian Lee. 2005. <a href="#">Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales</a> . In <i>Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)</i> , pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.	941
		942
		943
		944
		945
		946
		947
	Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. <a href="#">Language models as knowledge bases?</a> In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.	948
		949
		950
		951
		952
		953
		954
		955
		956
	Lihua Qian, Lin Qiu, Weinan Zhang, Xin Jiang, and Yong Yu. 2019. <a href="#">Exploring diverse expressions for paraphrase generation</a> . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 3173–3182, Hong Kong, China. Association for Computational Linguistics.	957
		958
		959
		960
		961
		962
		963
		964
	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019a. <a href="#">Language models are unsupervised multitask learners</a> .	965
		966
		967
	Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019b. <a href="#">Language models are unsupervised multitask learners</a> . <i>OpenAI blog</i> , 1(8):9.	968
		969
		970
		971
	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. <a href="#">Exploring the limits of transfer learning with a unified text-to-text transformer</a> . <i>CoRR</i> , abs/1910.10683.	972
		973
		974
		975
		976
	Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. 2019. <a href="#">On the convergence of adam and beyond</a> . <i>CoRR</i> , abs/1904.09237.	977
		978
		979

980	Steven J. Rennie, Etienne Marcheret, Youssef Mroueh,	Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa,	1037
981	Jerret Ross, and Vaibhava Goel. 2016. <a href="#">Self-critical</a>	Isabel Kloumann, Artem Korenev, Punit Singh Koura,	1038
982	<a href="#">sequence training for image captioning</a> . <i>CoRR</i> ,	Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Di-	1039
983	<a href="#">abs/1612.00563</a> .	ana Liskovich, Yinghai Lu, Yuning Mao, Xavier Mar-	1040
984	Abhilasha Sancheti, Balaji Vasan Srinivasan, and	tinet, Todor Mihaylov, Pushkar Mishra, Igor Moly-	1041
985	Rachel Rudinger. 2022. <a href="#">Entailment relation aware</a>	bog, Yixin Nie, Andrew Poulton, Jeremy Reizen-	1042
986	<a href="#">paraphrase generation</a> . <i>Proceedings of the AAAI</i>	stein, Rashi Rungta, Kalyan Saladi, Alan Schelten,	1043
987	<i>Conference on Artificial Intelligence</i> , 36(10):11258–	Ruan Silva, Eric Michael Smith, Ranjan Subrama-	1044
988	11266.	nian, Xiaoqing Ellen Tan, Binh Tang, Ross Tay-	1045
989	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec	lor, Adina Williams, Jian Xiang Kuan, Puxin Xu,	1046
990	Radford, and Oleg Klimov. 2017. <a href="#">Proximal policy</a>	Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan,	1047
991	<a href="#">optimization algorithms</a> . <i>CoRR</i> , <a href="#">abs/1707.06347</a> .	Melanie Kambadur, Sharan Narang, Aurelien Rod-	1048
992	Weijia Shi, Xiaochuang Han, Hila Gonen, Ari Holtzman,	riguez, Robert Stojnic, Sergey Edunov, and Thomas	1049
993	Yulia Tsvetkov, and Luke Zettlemoyer. 2022. <a href="#">Toward</a>	Scialom. 2023. <a href="#">Llama 2: Open foundation and fine-</a>	1050
994	<a href="#">human readable prompt tuning: Kubrick’s the shining</a>	<a href="#">tuned chat models</a> .	1051
995	<a href="#">is a good movie, and a good prompt too?</a>	Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan	1052
996	Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric	Kobyzev, and Ali Ghodsi. 2023. <a href="#">DyLoRA:</a>	1053
997	Wallace, and Sameer Singh. 2020. <a href="#">AutoPrompt: Elic-</a>	<a href="#">Parameter-efficient tuning of pre-trained models us-</a>	1054
998	<a href="#">iting Knowledge from Language Models with Auto-</a>	<a href="#">ing dynamic search-free low-rank adaptation</a> . In <i>Pro-</i>	1055
999	<a href="#">matically Generated Prompts</a> . In <i>Proceedings of the</i>	<i>ceedings of the 17th Conference of the European</i>	1056
1000	<i>2020 Conference on Empirical Methods in Natural</i>	<i>Chapter of the Association for Computational Lin-</i>	1057
1001	<i>Language Processing (EMNLP)</i> , pages 4222–4235,	<i>guistics</i> , pages 3274–3287, Dubrovnik, Croatia. As-	1058
1002	Online. Association for Computational Linguistics.	sociation for Computational Linguistics.	1059
1003	A. B. Siddique, Samet Oymak, and Vagelis Hristidis.	Ashwin Vijayakumar, Michael Cogswell, Ramprasaath	1060
1004	2020. <a href="#">Unsupervised paraphrasing via deep rein-</a>	Selvaraju, Qing Sun, Stefan Lee, David Crandall,	1061
1005	<a href="#">forcement learning</a> . In <i>Proceedings of the 26th</i>	and Dhruv Batra. 2018. <a href="#">Diverse beam search for</a>	1062
1006	<i>ACM SIGKDD International Conference on Knowl-</i>	<a href="#">improved description of complex scenes</a> . <i>Proceed-</i>	1063
1007	<i>edge Discovery &amp; Data Mining, KDD ’20</i> , page	<i>ings of the AAAI Conference on Artificial Intelligence</i> ,	1064
1008	1800–1809, New York, NY, USA. Association for	32(1).	1065
1009	Computing Machinery.	Maxim Kuznetsov Vladimir Vorobev. 2023. <a href="#">Chatgpt</a>	1066
1010	Richard Socher, Alex Perelygin, Jean Wu, Jason	<a href="#">paraphrases dataset</a> .	1067
1011	Chuang, Christopher D. Manning, Andrew Ng, and	Ellen M. Voorhees and Dawn M. Tice. 2000. <a href="#">Building</a>	1068
1012	Christopher Potts. 2013. <a href="#">Recursive deep models for</a>	<a href="#">a question answering test collection</a> . In <i>Proceedings</i>	1069
1013	<a href="#">semantic compositionality over a sentiment treebank</a> .	<i>of the 23rd Annual International ACM SIGIR Confer-</i>	1070
1014	In <i>Proceedings of the 2013 Conference on Empiri-</i>	<i>ence on Research and Development in Information</i>	1071
1015	<i>cal Methods in Natural Language Processing</i> , pages	<i>Retrieval, SIGIR ’00</i> , page 200–207, New York, NY,	1072
1016	1631–1642, Seattle, Washington, USA. Association	USA. Association for Computing Machinery.	1073
1017	for Computational Linguistics.	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	1074
1018	Tianxiang Sun, Yunfan Shao, Hong Qian, Xuan-	Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. 2022.	1075
1019	jing Huang, and Xipeng Qiu. 2022. <a href="#">Black-box</a>	<a href="#">Chain of thought prompting elicits reasoning in large</a>	1076
1020	<a href="#">tuning for language-model-as-a-service</a> . <i>CoRR</i> ,	<a href="#">language models</a> . <i>CoRR</i> , <a href="#">abs/2201.11903</a> .	1077
1021	<a href="#">abs/2201.03514</a> .	Jason Wei and Kai Zou. 2019. <a href="#">EDA: Easy data augmen-</a>	1078
1022	Richard S. Sutton, David McAllester, Satinder Singh,	<a href="#">tation techniques for boosting performance on text</a>	1079
1023	and Yishay Mansour. 1999. Policy gradient methods	<a href="#">classification tasks</a> . In <i>Proceedings of the 2019 Con-</i>	1080
1024	for reinforcement learning with function approxima-	<i>ference on Empirical Methods in Natural Language</i>	1081
1025	tion. In <i>Proceedings of the 12th International Con-</i>	<i>Processing and the 9th International Joint Confer-</i>	1082
1026	<i>ference on Neural Information Processing Systems</i> ,	<i>ence on Natural Language Processing (EMNLP-</i>	1083
1027	NIPS’99, page 1057–1063, Cambridge, MA, USA.	<i>IJCNLP)</i> , pages 6382–6388, Hong Kong, China. As-	1084
1028	MIT Press.	sociation for Computational Linguistics.	1085
1029	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	Qingru Zhang, Minshuo Chen, Alexander Bukharin,	1086
1030	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	Pengcheng He, Yu Cheng, Weizhu Chen, and	1087
1031	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	Tuo Zhao. 2023. <a href="#">Adaptive budget allocation for</a>	1088
1032	Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton	<a href="#">parameter-efficient fine-tuning</a> .	1089
1033	Ferrer, Moya Chen, Guillem Cucurull, David Esiobu,	Susan Zhang, Stephen Roller, Naman Goyal, Mikel	1090
1034	Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,	Artetxe, Moya Chen, Shuohui Chen, Christopher De-	1091
1035	Cynthia Gao, Vedanuj Goswami, Naman Goyal, An-	wan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mi-	1092
1036	thony Hartshorn, Saghar Hosseini, Rui Hou, Hakan	haylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel	1093

1094 Simig, Punit Singh Koura, Anjali Sridhar, Tianlu 1142  
1095 Wang, and Luke Zettlemoyer. 2022a. *Opt: Open* 1143  
1096 *pre-trained transformer language models.* 1144

1097 Tianjun Zhang, Xuezhi Wang, Denny Zhou, Dale Schu- 1145  
1098 urmans, and Joseph E. Gonzalez. 2022b. *Tempera:* 1146  
1099 *Test-time prompting via reinforcement learning.* 1147

1100 Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. 1148  
1101 Weinberger, and Yoav Artzi. 2020. *Bertscore: Eval-* 1149  
1102 *uating text generation with bert.* In *International* 1150  
1103 *Conference on Learning Representations.* 1151

1104 Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. 1152  
1105 *Character-level convolutional networks for text clas-* 1153  
1106 *sification.* *CoRR*, abs/1509.01626. 1154

1107 Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex 1155  
1108 Smola. 2022c. *Automatic chain of thought prompt-* 1156  
1109 *ing in large language models.* 1157

1110 Jianing Zhou and Suma Bhat. 2021. *Paraphrase genera-* 1158  
1111 *tion: A survey of the state of the art.* In *Proceedings* 1159  
1112 *of the 2021 Conference on Empirical Methods in Nat-* 1160  
1113 *ural Language Processing*, pages 5075–5086, Online 1161  
1114 and Punta Cana, Dominican Republic. Association 1162  
1115 for Computational Linguistics. 1163

1116 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, 1164  
1117 Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy 1165  
1118 Ba. 2023. *Large language models are human-level* 1166  
1119 *prompt engineers.* 1167

1120 **A Baseline LM Tuning Techniques**

1121 **Gradient-Search (GS):** The *GS* technique is based 1168  
1122 on the recent AutoPrompt (Shin et al., 2020) 1169  
1123 method, which optimizes task instructions with- 1170  
1124 out updating any parameters in the model. The 1171  
1125 search process begins in the vocabulary space, op- 1172  
1126 timizing the change in label log-likelihood when 1173  
1127 replacing token  $p_i$  in the task instruction with an- 1174  
1128 other token  $v$  from the vocabulary set. In our imple- 1175  
1129 mentation, each search iteration randomly selects 1176  
1130 one mini-batch of training examples and then ran- 1177  
1131 domly selects a token from the task instruction 1178  
1132 to update. The top  $k$  candidate tokens are deter- 1179  
1133 mined based on the approximate change in label 1180  
1134 log-likelihood:  $Top_v \{w_v^T \cdot \nabla_{w_{p_i}} \log P_{lm}(y|p, x)\}$ , 1181  
1135 where  $w_v$  is the embedding vector of a candidate 1182  
1136 token  $v$ . The resulting  $k$  new task instructions are 1183  
1137 evaluated again using label log-likelihood on the 1184  
1138 same training examples<sup>3</sup>, and the top-performing 1185  
1139 instruction is retained for the next search iteration. 1186  
1140 Prompt optimization always uses the original input 1187  
1141  $x$  when searching new task prompts (Shin et al., 1188

<sup>3</sup>The original AutoPrompt evaluates the new candidate instructions on another training mini-batch. For fewshot classification, we re-use the drawn training mini-batch to evaluate the complete new candidate instructions.

2020; Deng et al., 2022). In our work, we investi- 1142  
1143 gate the impact of incorporating paraphrases of  $x$  1144  
1145 during search. 1146

**Input-Finetuning (InTune):** As a straightfor- 1147  
1148 ward and efficient tuning technique, we compare 1149  
1150 to updating only the input embedding table in the 1151  
1152 transformer architecture. This method requires gra- 1153  
1154 dient computation similar to All-Finetuning (*All-* 1154  
1155 *Tune*) as well as the *GS* method. 1156

**LM-Head-Finetuning (HTune):** The 1157  
1158 transformer-based pre-trained language models 1158  
1159 consist of a language modeling head, which maps 1159  
1160 the hidden vectors to the token logit for each 1160  
1161 token in the vocabulary. For the *HTune* technique, 1161  
1162 we solely update the parameters of the language 1162  
1163 modeling head. 1163

**Classifier-Finetuning (ClsTune):** In *ClsTune*, 1164  
1165 we first create a feature representation  $h(x)$  for the 1164  
1166 input text  $x$  using average pooling of the final hid- 1165  
1167 den vectors in the last layer of the language model. 1166  
1168 Here, we assume that the language model (feature 1167  
1169 extractor) remains fixed, and we then construct a 1168  
1170 two-layer feedforward network with the *gelu* ac- 1169  
1171 tivation function (Hendrycks and Gimpel, 2016) 1170  
1172 as a classification module on top of the language 1171  
1173 model. 1172

**Softprompt-Tuning (SpTune):** In *SpTune* 1173  
1174 (Lester et al., 2021),  $L$  prompt tokens are 1174  
1175 prepended to the task instruction. These  $L$  tokens 1175  
1176 are associated with  $L$  dedicated prompt embedding 1176  
1177 vectors, extending the sequence of vectors derived 1177  
1178 from the task instruction and input text with an 1178  
1179 additional  $L$  trainable feature vectors. During train- 1179  
1180 ing, the original embedding table of the transformer 1180  
1181 model remains fixed, while a new prompt embed- 1181  
1182 ding table is trained by backpropagating the label 1182  
1183 log-likelihood into the prompt embedding table. In 1183  
1184 contrast to *InTune*, here the prompt vectors do not 1184  
1185 need to map to vocabulary words. 1185

**Low-Rank Adaptation (LoRA):** *LoRA* is one 1186  
1187 of the latest efficient-tuning techniques specifically 1186  
1188 designed for PLMs (Hu et al., 2021). It learns 1187  
1189 low-rank adaptation matrices for the query and 1188  
1190 value weight matrices within the transformer model. 1189  
1191 For a pre-trained weight matrix  $W_q \in \mathbb{R}^{d \times k}$ , 1190  
1192 *LoRA* learns the necessary adaptation (i.e., mod- 1191  
1193 ification) of the weight matrix for a downstream 1192  
1194 task through a low-rank decomposition, expressed 1193  
1195 as  $W_q + \Delta W_q \approx W_q + BA$ . Here,  $B \in \mathbb{R}^{d \times r}$ , 1194  
1196  $A \in \mathbb{R}^{r \times k}$ , and the rank  $r \leq \min(d, k)$ . The 1195  
1197 adaptation matrices  $A$  and  $B$  are the only param- 1196  
1198 eters subject to training, while the original matrix 1197  
1199 1198  
1200 1199

1194  $W_q$  does not receive any gradient updates. Studies  
 1195 have shown that *LoRA* performs on par with, or  
 1196 better than, *AllTune* across various PLMs (Hu et al.,  
 1197 2021).  
 1198 All language model tuning techniques we have dis-  
 1199 cussed will use the same input format. For example  
 1200 in the sentiment classification task, we use the fol-  
 1201 lowing format:

1202 “< s > {instruction} {text} . It was < mask > . < / s > ”.

1203 Except for *Clstune*, all of our tuning techniques  
 1204 maximize the probability of the correct label token  
 1205 in place of the < mask > token. In contrast, *Clstune*  
 1206 takes the formatted input and classifies it into one  
 1207 of the predefined class labels.

## 1208 B Few-shot Paraphrase Fine-Tuning 1209 (Further Results)

1210 This section provides additional results that com-  
 1211 pare our training objectives for fine-tuning the para-  
 1212 phrase generator using the feedback from the down-  
 1213 stream language model.

1214 The off-policy learning technique improves per-  
 1215 formance when using basic rewards (i.e., 69.1%  
 1216 compared to 67.9% with mixed decoding). How-  
 1217 ever, the combined effect of off-policy learning  
 1218 and reward normalization decreases performance.  
 1219 With mixed decoding, ‘PG-Z’ yields an accuracy of  
 1220 71.2% in on-policy learning compared to an accu-  
 1221 racy of 68.0% with off-policy learning. The ‘AVG’  
 1222 column in Table 4 further verifies this conclusion  
 1223 that reward normalization is not improving the fi-  
 1224 nal performance while training the model with off-  
 1225 policy learning. We hypothesize that with the off-  
 1226 policy learning technique, the normalized rewards  
 1227 should be re-weighted properly if the sampled para-  
 1228 phrases are from the fixed paraphrase model.

## 1229 C Example Paraphrases

1230 In Table 5 and Table 7, we present example para-  
 1231 phrases from the AGNews test dataset generated by  
 1232 our pre-trained paraphrase model. We selected the  
 1233 AGNews dataset for its suitability in paraphrasing  
 1234 longer texts or short paragraphs. Subsequently, in  
 1235 Table 6 and Table 8, we display the generated para-  
 1236 phrases after fine-tuning the paraphrase model with  
 1237 the RIFF objective. Fewer hallucinations can be  
 1238 observed in the new paraphrases, which are high-  
 1239 lighted in red.

## D Further Training Details 1240

1241 The learning rate for each LM tuning technique  
 1242 was separately fine-tuned from the set {0.5, 0.3,  
 1243 0.1, 0.01, 0.001, 0.0001, 0.00001} using the  
 1244 train/validation split created for the seed 11 on  
 1245 the SST2 dataset. The tuned learning rates were  
 1246 then applied globally across other datasets and ex-  
 1247 periments. For paraphrase fine-tuning, we train  
 1248 all the parameters in T5-base with the learning  
 1249 rate of 0.00001. In Tables 9 and 10, we list the  
 1250 hyper-parameters and learning rates used across  
 1251 all datasets. For optimization, we utilized the  
 1252 AdamW (Loshchilov and Hutter, 2017)<sup>4</sup> optimizer  
 1253 with the AMSGrad variant set to True (Reddi et al.,  
 1254 2019). We implemented the methods using the  
 1255 HuggingFace<sup>5</sup> library and the PyTorch<sup>6</sup> machine  
 1256 learning framework. We report the accuracy met-  
 1257 ric on these classification datasets. The experi-  
 1258 ments were conducted using multiple NVIDIA’s  
 1259 A40 GPU cards.

## E Paraphrases for Few-shot LM Tuning (Further Results) 1260

1261 Due to space limitations, we present the results for  
 1262 *Clstune*, *HTune*, and *InTune* in Table 11. 1263

## F Extended Related Works 1264

1265 **Prompt Optimization & Efficient Tuning:** Re-  
 1266 cent research proposes various techniques for  
 1267 prompt optimization and efficient tuning of lan-  
 1268 guage models. In our experiments, we have used  
 1269 successful techniques from each of these areas.

1270 FluentPrompt (Shi et al., 2022) is a recent dis-  
 1271 crete prompting technique based on the projected  
 1272 gradient-descent and Langevin dynamics. Flu-  
 1273 entPrompt introduces a fluency constraint within  
 1274 Langevin dynamics to generate a sample of high-  
 1275 performing prompts for more interpretable analysis  
 1276 of these discrete prompts. The optimized prompts  
 1277 by FluentPrompt performs on-par to the Auto-  
 1278 Prompt, however they have lower perplexity (Shi  
 1279 et al., 2022).

1280 Building upon *SpTune* (Lester et al., 2021) and  
 1281 P-tuning (Li and Liang, 2021), P-tuning V2 (Liu  
 1282 et al., 2022) introduced the concept of deep prompt  
 1283 tuning. This method involves injecting prompt  
 1284 vectors into the deeper layers of the transformer

<sup>4</sup><https://pytorch.org/docs/stable/generated/torch.optim.AdamW.html>

<sup>5</sup><https://huggingface.co/>

<sup>6</sup><https://pytorch.org/>

Table 4: The accuracy of the best performing validation checkpoint in the 128-shot SST2 classification task trained with the off-policy learning technique.

Learn Tech	Off-Policy			AVG
	Top-P	Beam	Mixed	
No Tuning	67.5	67.5	67.5	67.5
PG	68.6 (1.8   67.5)	68.4 (1.6   67.4)	69.1 (1.6   67.5)	68.7 (1.7   67.5)
PG-Z	68.8 (1.7   67.5)	68.7 (1.1   67.4)	68.0 (2.1   67.2)	68.5 (1.6   67.4)
MML	<b>69.2</b> (2.8   68.0)	<b>70.1</b> (2.4   68.6)	<b>70.1</b> (3.3   68.4)	<b>69.8</b> (2.7   68.3)
MML-Z	<b>69.2</b> (2.5   68.3)	69.7 (3.5   68.6)	<b>70.1</b> (2.8   68.6)	69.7 (2.9   68.5)

Table 5: Generated top eight paraphrases using diverse beam search **with the pre-trained paraphrase model (distilled from GPT3.5-turbo)** for the following input sentence from the AGNews test dataset:

“Gardner Loses Quest for Repeat Wrestling Gold US heavyweight Rulon Gardner lost his Olympic title Wednesday after being beaten in the semi-final stage of the 120kg Greco-Roman wrestling event by Georgiy Tsurtsumia of Kazakhstan.” Contradicting information or hallucinations are highlighted in **red**.

Paraphrase	Text
P1	“On Wednesday, US heavyweight Rulon Gardner lost his Olympic title after losing in the semi-finals of the 120kg Greco-Roman wrestling event against Georgiy Tsurtsumia of Kazakhstan.”
P2	“Following his loss to Georgiy Tsurtsumia of Kazakhstan in the semi-finals of the 120kg Greco-Roman wrestling event, US heavyweight Rulon Gardner lost his bid for repeat Wrestling Gold.”
P3	“Rulon Gardner, the heavyweight world champion, suffered a major upset on Wednesday as he was defeated in the semi-finals of the 120kg Greco-Roman wrestling tournament by Georgiy Tsurtsumia of Kazakhstan. Gardner lost his Olympic title bid to <b>George Kubica after losing 1-1 and 6-1.</b> ”
P4	“US heavyweight Rulon Gardner suffered a setback in the <b>final round</b> of his Olympic title against Georgiy Tsurtsumia of Kazakhstan, <b>leading to an upset loss that could have resulted in injury.</b> ”
P5	“Despite winning the Olympic title <b>in 2004</b> , Rulon Gardner, US heavyweight champion, lost his quest to repeat as a wrestler in the 120kg Greco-Roman wrestling event after losing to Georgiy Tsurtsumia of Kazakhstan in their semi-final match.”
P6	“The US heavyweight champion, Rulon Gardner of the United States, suffered a defeat to Georgiy Tsurtsumia from Kazakhstan in the semi-finals of his men’s 120kg Greco-Roman wrestling tournament, which cost him the Olympic championship on Wednesday.”
P7	“In a bid to avoid repeat Wrestling Gold, US heavyweight Rulon Gardner lost his Olympic title on Wednesday after losing in the <b>final round of Georges Tsurtsumia’s</b> 120kg Greco-Roman wrestling match against Kazakhstan’s <b>Rorum Ignite.</b> ”
P8	“ <b>After</b> losing his Olympic title on Wednesday, US heavyweight Rulon Gardner was defeated in the semi-finals of Greco-Roman wrestling event at 120kg against Georgiy Tsurtsumoj, Kazakhstan.”

Table 6: Generated top eight paraphrases using diverse beam search **after fine-tuning the paraphrase model with the RIFF method** in the 128-shot setting. The example input sentence is from the AGNews test dataset:

“Gardner Loses Quest for Repeat Wrestling Gold US heavyweight Rulon Gardner lost his Olympic title Wednesday after being beaten in the semi-final stage of the 120kg Greco-Roman wrestling event by Georgiy Tsurtsumia of Kazakhstan.” Contradicting information or hallucinations are highlighted in **red**

Paraphrase	Text
P1	“The US heavyweight Rulon Gardner lost his Olympic title on Wednesday after losing in the semi-final stage of the 120kg Greco-Roman wrestling event to Georgiy Tsurtsumia of Kazakhstan.”
P2	“On Wednesday, Rulon Gardner lost his Olympic title after losing in the semi-final stage of the 120kg Greco-Roman wrestling event to Georgiy Tsurtsumia of Kazakhstan.”
P3	“Georgiy Tsurtsumia of Kazakhstan defeated US heavyweight Rulon Gardner in the semi-final stage of the 120kg Greco-Roman wrestling event, resulting in Gardner losing his Olympic title on Wednesday.”
P4	“ <b>Despite</b> losing his Olympic title on Wednesday, US heavyweight Rulon Gardner lost his quest for <b>repeat wrestling</b> . He was defeated in the 120kg Greco-Roman wrestling event by Georgiy Tsurtsumia of Kazakhstan in their semi-final stage.”
P5	“Rulon Gardner, the heavyweight US heavy weight loser of Olympic title after losing in the semi-final stage of the 120kg Greco-Roman wrestling event to Georgiy Tsurtsumia of Kazakhstan.”
P6	“US heavyweight Rulon Gardner suffered a loss in the semi-final stage of the 120kg Greco-Roman wrestling event <b>after</b> losing his Olympic title.”
P7	“In the 120kg Greco-Roman wrestling event, Rulon Gardner lost his Olympic title on Wednesday after losing in the semi-final stage by Georgiy Tsurtsuma of Kazakhstan.”
P8	“ <b>After</b> losing his Olympic title on Wednesday, Rulon Gardner lost the semi-final stage of a 120kg Greco-Roman wrestling event against Georgiy Tsurtsumoia of Kazakhstan in the <b>Quest for Repeat Wrestling.</b> ”

Table 7: Generated top eight paraphrases using diverse beam search **with the pre-trained paraphrase model (distilled from GPT3.5-turbo)** for the following input sentence from the AGNews test dataset:

“Calif. Aims to Limit Farm-Related Smog (AP) AP - Southern California’s smog-fighting agency went after emissions of the bovine variety Friday, adopting the nation’s first rules to reduce air pollution from dairy cow manure.”  
 Contradicting information or hallucinations are highlighted in **red**.

Paraphrase	Text
P1	“The smog-fighting agency in southern California followed <b>the lead of other states</b> by implementing the first rules to control air pollution from dairy cow manure emissions.”
P2	“Southern California’s smog-fighting agency took action on bovine eutrophication by adopting the first rules of the nation to reduce farm-related air pollution, including dairy cow manure emissions.”
P3	“California’s smog-fighting agency in southern California took action on bovine deposition Friday, following the implementation of the first rules in the country to control air pollution from dairy cow manure.”
P4	“Calf Air Pollution Control: Southern California’s smog-fighting agency took action on Friday to limit farm-related <b>sewage</b> emissions by adopting the first rules in the country to reduce air pollution from dairy cow manure.”
P5	“Efforts to Reduce Farm-Related Smog: Calif.’s agency in Southern California set the pace on Friday with its first-ever goals to combat bovine emissions and the <b>nation’ll be moving quickly</b> to eliminate farm-related air pollution, <b>including dairy</b> cow manure.”
P6	“Southern California’s <b>EPA</b> sought to reduce air pollution from farm-related manure, and the agency adopted the first rules <b>in history</b> to address bovine smog emissions.”
P7	“Calif. aims to reduce farm-related smog, which Southern California’s <b>first specialized</b> agency for reducing bovine emissions took action on Friday by adopting the <b>nation’ reviews</b> that address this issue with “‘first ever’” rules from the <b>U.S. federal government</b> .”
P8	“The smog-fighting agency in Southern California adopted new regulations on Friday to reduce farm-related pollution, including air pollution from dairy cow manure.”

Table 8: Generated top eight paraphrases using diverse beam search **after fine-tuning the paraphrase model with the RIFF method** in the 128-shot setting. The example input sentence is from the AGNews test dataset:

“Calif. Aims to Limit Farm-Related Smog (AP) AP - Southern California’s smog-fighting agency went after emissions of the bovine variety Friday, adopting the nation’s first rules to reduce air pollution from dairy cow manure.”  
 Contradicting information or hallucinations are highlighted in **red**

Paraphrase	Text
P1	“The smog-fighting agency in southern California adopted the first rules to reduce air pollution from dairy cow manure, targeting bovine emissions.”
P2	“Southern California’s smog-fighting agency took action on bovine emissions by adopting the nation’ first rules to reduce air pollution from dairy cow manure.”
P3	“On Friday, Southern California’s smog-fighting agency adopted the nation’ first rules to reduce air pollution from dairy cow manure emissions.”
P4	“California’s smog-fighting agency took action on bovine emissions by adopting the nation’ first rules to reduce air pollution from dairy cow manure.”
P5	“Calif. seeks to reduce farm-related smog emissions by adopting the nation’s first rules on Friday, which targets bovine emission in Southern California and is <b>one of the country’ best agencies</b> for curbing air pollution from dairy cow manure.”
P6	“Southern California’s scounding agency took action on bovine emissions by adopting the nation’ first rules to reduce air pollution from dairy cow manure.”
P7	“The smog-fighting agency of Southern California targeted bovine emission on Friday, adoptting the country’s first rules to reduce air pollution from dairy cow manure.”
P8	“AP: Southern California’s <b>federal</b> agency adopting the nation’ first rules to reduce air pollution from dairy cow manure.”



Table 9: Shared hyper-parameters used across all experiments and datasets.

Hyper-parameter	Value
Top- $k$ candidates in <i>GS</i>	$k=4$
batch size (RoBERTa-large)	8
batch size in <i>GS</i> (RoBERTa-large)	2
Weight decay	0.0001
Max epochs	100
length cutoff	128 tokens
Paraphrase sample size	$M=8$
Checkpointing steps	8
$D'$ in <i>ClsTune</i>	128
Prompt len in <i>SpTune</i>	$L=25$
$\beta$ in MML	0.1
$\beta$ in PG	0.6
<i>LoRA</i> $\alpha$	32
<i>LoRA</i> $r$	8
<i>LoRA</i> dropout	0.1
Diversity penalty for Div beam	3.0
Repetition penalty for Div beam	10.0
Temperature in Div beam	0.7
P value for top-p	0.99

Table 10: Learning rates used per Language Model (LM) tuning technique.

LM Tuning Technique	Learning Rate
<i>GS</i>	No rate
<i>AllTune</i>	0.00001
<i>InTune</i>	0.001
<i>HTune</i>	0.001
<i>ClsTune</i>	0.001
<i>SpTune</i>	0.001
<i>LoRA</i>	0.0001

model to close the performance gap with AllTuning in medium-sized language models. We have experimented with *LoRA* (Hu et al., 2021), a recent low-rank adaptation technique for tuning language models. Other potential methods include training bottleneck adapter modules (Houlsby et al., 2019b; Lin et al., 2020) added per sub-layer of the transformer model. *LoRA* outperforms adapter tuning and P-Tuning V2 techniques (Hu et al., 2021). The successors of *LoRA* include DyLoRA (Valipour et al., 2023) which dynamically learns a range of adaptation ranks, thus eliminating the need to search the rank of the adaptation matrices as a hyper-parameter. Similarly, *AdaLoRA* dynamically allocates the parameter budget among the weight matrices during adaptation, with matrices of higher

priority (i.e., those with greater importance to the downstream task) receiving higher adaptation ranks than less important matrices (Zhang et al., 2023).

In scenarios where gradients are absent, Black-Box Tuning (Sun et al., 2022) applies derivative-free algorithms for optimizing continuous prompts. For discrete prompt optimization, RLPrompt (Deng et al., 2022) employs the on-policy version of soft Q-learning (Guo et al., 2021) to find the optimal prompt tokens in a gradient-free setting. Decoder Tuning (Cui et al., 2023) learns a decoder network over the language model, thus circumventing the need for gradient computation and input-side prompt tuning in few-shot classification. In a recent study, TEMPERA (Zhang et al., 2022b) introduced a novel approach that involves test-time discrete prompt editing using a trained RL agent. This agent is capable of modifying the instruction, in-context examples, or the verbalizers based on the given task input.

The use of Language Models (LLMs) in generating instructions for downstream tasks has involved a two-step process. Initially, LLMs generate a set of candidate instructions, and subsequently, the highest-scoring instruction is utilized to prompt another LLM to perform the downstream task. This approach, known as prompt-based generation-then-filtering, has been investigated in the recent APE method (Zhou et al., 2023). APE demonstrates the ability to generate prompts that achieve performance comparable to human-designed prompts (Zhou et al., 2023).

To prompt language models for reasoning tasks, another line of research augment the input context with demonstration examples outlining the intermediate reasoning steps to form the answer. Providing manually or automatically generated chain-of-thoughts within these demonstrations strikingly improve LLMs performance in reasoning tasks (Wei et al., 2022; Zhang et al., 2022c; Kojima et al., 2022).

All of the aforementioned techniques for prompt optimization and efficient tuning of the language model use the original task’s input text (or the original input context) provided within the dataset.

**RL for Paraphrase Generation:** In the following paragraphs, we provide a brief overview of similar reinforcement learning objectives employed for paraphrase generation. Li et al. (Li et al., 2018) used a deep RL technique, training a pointer-generator network as the paraphrase generator and a decomposable attention model as the evaluator

Table 11: Average accuracy on the standard evaluation sets for the 16-shot text classification using the *ClsTune*, *HTune*, and *InTune* fine-tuning techniques. Numbers in parentheses are standard deviations across the five train/validation folds. The last column is the micro averaged performance across the datasets.

Tuning Method	SST2	SST5	CR	MR	TREC	AGN	AVG
<i>ClsTune</i>	72.6 (2.4)	34.4 (2.6)	71.4 (2.7)	67.3 (2.8)	74.8 (3.4)	81.7 (1.6)	70.9 (2.2)
+RIFF (train)	72.5 (3.4)	33.9 (3.7)	68.3 (3.9)	70.3 (0.9)	75.8 (1.7)	84.0 (0.9)	71.9 (2.0)
+RIFF (train+test)	74.0 (3.3)	35.0 (3.6)	71.1 (4.4)	72.0 (1.7)	76.8 (2.9)	84.9 (0.9)	73.3 (2.1)
<i>HTune</i>	87.4 (2.3)	37.4 (2.0)	84.0 (2.8)	83.1 (1.8)	62.4 (7.4)	81.4 (1.4)	76.0 (2.0)
+RIFF (train)	88.1 (1.7)	40.3 (1.9)	84.5 (1.5)	83.4 (2.8)	70.7 (4.8)	83.4 (0.9)	77.8 (1.6)
+RIFF (train+test)	89.1 (1.2)	40.4 (1.8)	86.4 (0.8)	83.1 (3.7)	71.6 (5.9)	85.2 (1.1)	79.0 (1.6)
<i>InTune</i>	91.5 (1.2)	42.3 (4.2)	87.3 (2.0)	84.0 (2.5)	67.7 (5.8)	83.8 (2.2)	78.9 (2.5)
+RIFF (train)	92.6 (0.3)	43.2 (1.9)	87.5 (1.8)	85.9 (2.3)	63.8 (5.3)	85.6 (0.8)	80.2 (1.3)
+RIFF (train+test)	93.1 (0.6)	43.9 (2.3)	89.0 (1.8)	86.0 (2.4)	69.6 (6.3)	86.9 (0.4)	81.3 (1.3)

1353 which assigns a paraphrase score to pairs of sentences. The generator was trained using the policy gradient objective, with reward shaping and scaling to stabilize the training process (Li et al., 2018). Another approach by Qian et al. (Qian et al., 2019) focused on generating diverse paraphrases by training multiple generators, accompanied by a paraphrase discriminator and a generator discriminator. Policy gradient objective and self-critical learning (Rennie et al., 2016) were employed for training the generators, with the baseline reward used in the policy gradient objective being the reward obtained from the greedy-decoded sequence. Liu et al. (Liu et al., 2020) also applied the policy gradient objective with self-critical learning, incorporating multiple reward functions such as Rouge score with the reference paraphrase, negative Rouge score with the input sentence to encourage lexical variations, and semantic similarity score between the paraphrase and the input sentence to ensure semantic fidelity.

1374 Another study by Du and Ji (Du and Ji, 2019) compared the use of imitation learning algorithm DAGGER with policy gradient REINFORCE for paraphrase generation. The policy gradient objective has also been applied in generating paraphrases while considering multiple objectives for entailment relation-aware paraphrase generation (Sanchez et al., 2022). In the context of chatbot responses, a recent work studies unsupervised paraphrase generation with proximal policy optimization, aiming to maximize a combination of rewards such as textual entailment, semantic similarity, language fluency, and lexical dissimilarity (Garg et al., 2021). Similarly, the policy gradient objective has been employed to optimize multiple rewards, similar to previous work, for un-

1390 supervised paraphrase generation (Siddique et al., 2020). 1391

1392 While previous studies have applied RL techniques for paraphrase generation, we propose the use of MML gradients instead of policy gradients to train our paraphrase model. Our training objective fine-tunes the paraphrase model for a downstream classification task. Our paraphrase model has been distilled from a large language model. 1397

## 1398 G Task Instructions & Input Format 1399

1400 Table 12 provides a summary of the task instructions that we append before the inputs, as well as the class verbalizers for classifying the input text. The instructions and input templates are derived from prior work in prompt optimization (Deng et al., 2022). 1405

Table 12: Number of classes  $C$ , test set size  $T$ , the input format, and the instruction used per dataset. The label words are provided within the instructions.

Dataset	$C$	$T$	Input Format	Instruction
SST2	2	1821	"<s> {Instruction} {Text} . It was <mask> . </s>"	"In this task, you are given sentences from movie reviews. The task is to classify a sentence as 'great' if the sentiment of the sentence is positive or as 'terrible' if the sentiment of the sentence is negative."
SST5	5	2210	"<s> {Instruction} {Text} . It was <mask> . </s>"	"In this task, you are given sentences from movie reviews. Based on the given review, classify it to one of the five classes: (1) terrible, (2) bad, (3) okay, (4) good, and (5) great."
CR	2	2000	"<s> {Instruction} {Text} . It was <mask> . </s>"	"In this task, you are given sentences from customer reviews. The task is to classify a sentence as 'great' if the sentiment of the sentence is positive or as 'terrible' if the sentiment of the sentence is negative."
MR	2	2000	"<s> {Instruction} {Text} . It was <mask> . </s>"	"In this task, you are given sentences from movie reviews. The task is to classify a sentence as 'great' if the sentiment of the sentence is positive or as 'terrible' if the sentiment of the sentence is negative."
TREC	6	500	"<s> {Instruction} <mask>: {Text} . </s>"	"You are given a question. You need to detect which category better describes the question. Answer with 'Description', 'Entity', 'Expression', 'Human', 'Location', and 'Number'."
AG's News	4	7600	"<s> {Instruction} <mask> News: {Text} . </s>"	"In this task, you are given a news article. Your task is to classify the article to one out of the four topics 'World', 'Sports', 'Business', 'Tech' if the article's main topic is relevant to the world, sports, business, and technology, correspondingly. If you are not sure about the topic, choose the closest option."