

Advancing Parameter Efficiency in Fine-tuning via Representation Editing

Anonymous ACL submission

Abstract

Parameter Efficient Fine-Tuning (PEFT) has gained significant attention for its ability to achieve competitive results while updating only a small subset of trainable parameters. Despite the promising performance of current PEFT methods, they present challenges in hyperparameter selection, such as determining the rank of LoRA or Adapter, or specifying the length of soft prompts. In addressing these challenges, we propose a novel approach to fine-tuning neural models, termed Representation EDiting (RED), which scales and biases the representation produced at each layer. RED substantially reduces the number of trainable parameters by a factor of 25, 700 compared to full parameter fine-tuning, and by a factor of 32 compared to LoRA. Remarkably, RED achieves comparable or superior results to full parameter fine-tuning and other PEFT methods. Extensive experiments were conducted across models of varying architectures and scales, including RoBERTa, GPT-2, T5, and Llama-2, and the results demonstrate the efficiency and efficacy of RED, positioning it as a promising PEFT approach for large neural models.

1 Introduction

Pre-training on large-scale unlabeled datasets and then fine-tuning on task-specific datasets has led to significant improvements across various natural language processing (NLP) tasks and has emerged as the predominant training paradigm (Devlin et al., 2018; Raffel et al., 2020; Radford et al., 2018). However, performing full parameter fine-tuning for each task would be prohibitively expensive with the growing model scale (Brown et al., 2020). For example, BERT consists of up to 220 million parameters; T5 comprises up to 11 billion parameters and GPT-3 contains up to 175 billion parameters. In this context, how to efficiently and effectively adapt large models to particular downstream tasks is an intriguing research issue (He et al., 2021).

To address this issue, researchers have proposed three main lines of Parameter Efficient Fine-Tuning (PEFT) methods (Ding et al., 2022). Specifically, additional-based methods introduce extra trainable neural modules or parameters that do not exist in the original model (Houlsby et al., 2019; Karimi Mahabadi et al., 2021; Li and Liang, 2021a; Lester et al., 2021a). Specification-based methods specify certain parameters in the original model become trainable, while others are frozen (Zaken et al., 2021; Guo et al., 2020). Reparameterization-based methods reparameterize trainable parameters to a parameter-efficient form by transformation (Hu et al., 2021; Zhang et al., 2023a; Ding et al., 2023).

Among these PEFT methods, Low-Rank Adaptation (LoRA) is considered one of the most efficient methods at present and its efficacy has been empirically validated across diverse models of varying scales. Despite its excellent performance, it still requires a considerable amount of trainable parameters. According to Aghajanyan et al. (2020) and Kopiczko et al. (2023), the upper bound for intrinsic dimensions is much smaller than what is typically utilized in such methods. For instance, the d_{90} ¹ for RoBERTa base is reported to be 896. Still, when utilizing LoRA to fine-tune this model, the number of trainable parameters reaches 0.3M, suggesting that the parameter count could be reduced further.

In addition, although previous works (Mao et al., 2021; He et al., 2021; Ding et al., 2022) have attempted to design different lightweight module structures or insert these modules into different positions in the base model, these PEFT methods consider fine-tuning the model from the perspective of adjusting model weights, which leads to many inconveniences in the selection of hyperparameters, such as the rank of LoRA and Adapter, as well as

¹ d_{90} denotes the smallest number of trainable parameters as being 90% of the full training metric.

the length of Soft Prompt and Prefix.

Inspired by the idea of representation engineering (Zou et al., 2023) representation can be modified to steer model outputs toward specific concepts and change the model’s behavior. We hypothesize that we can also consider fine-tuning the model from the perspective of editing neural network representations, leading to our proposed **Representation EDiting** (RED) approach. Instead of focusing on neurons and their connections, we fine-tune the model by learning a group of “edit vectors” to directly edit the representations of each layer and freezing the base model parameters, as shown in figure 1 (b).

Moreover, RED is highly parameter efficient. Using Llama-2 7B as an example, we show that RED can still achieve very promising performance by adjusting only 0.26M parameters, which is 25,700 times less than full parameter fine-tuning, making it both storage- and compute-efficient.

The contribution of this study can be summarized as follows:

- We consider fine-tuning the model from a new perspective of directly modifying the model representation, which is different from the previous work that adjusted the model weight, and propose our PEFT method, **Representation EDiting** (RED).
- We conducted extensive experiments on models with different structures and scales, including RoBERTa, GPT-2, T5, and Llama-2, and validated the effectiveness of RED on a series of NLU and NLG tasks although it only requires a small number of trainable parameters, and is quite simple to implement.
- We perform the ablation study to better understand the individual components of RED and their effects on performance.

2 Related Work

Ding et al. (2022) categorize the PEFT methods into three groups according to the operations on the learnable parameters: addition-based, specification-based, and reparameterization-based methods.

Addition-based methods introduce additional components for training based on the foundation model. Specifically, Houlsby et al. (2019); Stickland and Murray (2019); Karimi Mahabadi et al. (2021) and Rücklé et al. (2020) inject learnable bottleneck neural modules to the transformer layers.

Brown et al. (2020) and Shin et al. (2020) found that by concatenating some discrete tokens before the input text, the performance of the model can be improved without updating parameters. However, manually designing prompts requires a lot of effort, and the optimization problem in discrete space is relatively more difficult. Therefore, the subsequent works (Lester et al., 2021b; Li and Liang, 2021b; Wu et al., 2023; Wang et al., 2023) replace these discrete tokens with continuous vectors in front of the embedding layer or various hidden layers, also known as soft prompts, and optimize them through simple gradient descent.

Specification-based methods do not introduce any new parameters in the model, and they sparsely select part of the foundation model parameters for adjustment and freeze other parameters. Among them, Lee et al. (2019) adjusts the model parameters of the last few layers of BERT and RoBERTa. BitFit (Ben-Zaken et al., 2021) fine-tunes the model by only optimizing the bias terms inside the model. Unlike both of these methods, which manually specify the parameters that need to be adjusted in the network, Guo et al. (2020) and Zhao et al. (2020) use the learnable mask to dynamically select the parameters that need to be adjusted.

Reparameterization-based methods transform the optimization process of trainable parameters into a low-dimensional subspace. LoRA (Hu et al., 2021) proposes to employ low-rank matrices to approximate the weight changes during fine-tuning. QLoRA (Dettmers et al., 2023) combines low-rank adaptation with model quantization to further reduce storage usage during the model fine-tuning process. AdaLoRA (Zhang et al., 2023b) proposes using SVD decomposition to approximate the changes in weights, which allocate more trainable parameters to more important weight matrices, resulting in better performance.

What’s more, IA3 (Liu et al., 2022) and VeRA (Kopiczko et al., 2023) also consider scaling vectors in their implementation. However, IA3 not only needs to adjust the key vectors and value vectors of the attention module, as well as the vectors of the projection matrix but also needs to introduce additional complex loss functions. VeRA still needs to introduce a randomly initialized LoRA matrix and adjust the vectors passing through the projection matrix under LoRA. Compared to them, RED is simpler in implementation, as it only needs to directly edit the model’s representations. Moreover, RED has been extensively experimented on

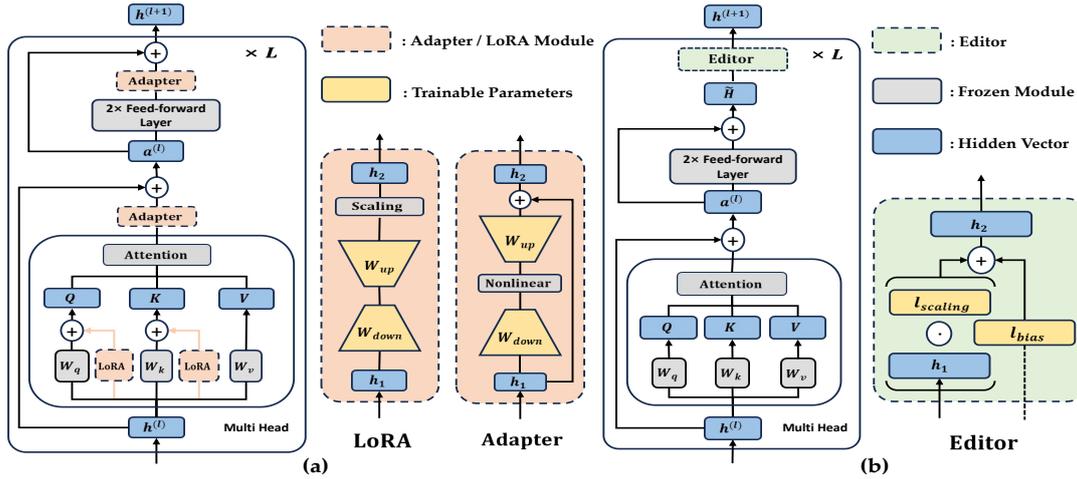


Figure 1: Comparison of previous PEFT methods and our proposed RED. *Left*: LoRA introduces the learnable bottleneck-shaped modules (orange area) through parallel connections for the \mathbf{W}_q and \mathbf{W}_v matrices of attention blocks, and models the weight changes of these two matrices in a low-rank manner. Adapter introduces learnable modules with similar structures (orange area) through serial connections after the attention sub-layers and feed-forward sub-layers to train the model. *Right*: RED introduces two learnable vectors $l_{scaling}$ and l_{bias} directly modifies the representation (green area) after feed-forward sub-layer to fine-tune the model.

various models of different architectures and scales, and its effectiveness has been fully demonstrated.

Representation engineering (Zou et al., 2023) suggests that neural representations are becoming more well-structured and place representations and transformations between them at the center of analysis rather than neurons or circuits. Specifically, Liu et al. (2023) points out that neural network weights determine neural activity, neural activity determines the networks’ output, and the networks’ output determines the networks’ behavior and utilizes this feature to operate in the representation space and achieves model alignment. Turner et al. (2023) adds a “steer vector” to the representation of each hidden layer during inference time to control the sentiment and style of the model output. Subramani et al. (2022) also extracted these “steel vectors” in the hidden space and completed unsupervised text style transfer by modifying the hidden representation through these vectors.

3 Method

In this section, we briefly review previous methods and introduce **Representation EDiting (RED)**, a novel parameter effective fine-tuning method that adapts pre-trained models to downstream tasks by directly modifying model representations.

3.1 Recap of previous PEFT methods

The transformer model (Vaswani et al., 2017) is now the cornerstone architecture behind most state-

of-the-art PLMs. Transformer models are composed of L stacked blocks, where each block contains two types of sub-layers: multi-head self-attention and fully connected feed-forward network (FFN). Except for the prompt-based methods which introduce learnable parameters in the embedding layer, many other PEFT methods are trained based on these two sub-layers.

Figure 1 (a) shows two commonly used PEFT methods, Adapter and LoRA. Except for a few additional parameters that need to be trained, the parameters of the pre-trained model are frozen.

Specifically, LoRA (Hu et al., 2021) introduces the learnable bottleneck-shaped modules through parallel connections for the \mathbf{W}_q and \mathbf{W}_v matrices of attention blocks and models the weight changes of these two matrices in a low-rank manner. For a pre-trained weight matrix $\mathbf{W} \in \mathbb{R}^{d \times k}$, LoRA represents its update with two low-rank decomposition matrices: $\mathbf{W} + \Delta\mathbf{W} = \mathbf{W} + s \cdot \mathbf{W}_{down} \mathbf{W}_{up}$, where $\mathbf{W}_{down} \in \mathbb{R}^{d \times r}$, $\mathbf{W}_{up} \in \mathbb{R}^{r \times k}$ and s is the scaling scalar, which is a hyperparameter set in advance. For $h = x\mathbf{W}$, LoRA modified forward pass yields:

$$h = x\mathbf{W} + s \cdot x\mathbf{W}_{down} \mathbf{W}_{up} \quad (1)$$

The initial adapter (Houlsby et al., 2019) inserts trainable adapter modules between transformer sub-layers. The adapter module contains a down-projection matrix $\mathbf{W}_{down} \in \mathbb{R}^{d \times r}$, map input

239 $h_1 \in \mathbb{R}^d$ to a low dimensional space of the spec- 289
 240 240
 241 241
 242 242
 243 243
 244 244
 245 245

$$246 \quad h_2 = h_1 + f(h_1 \mathbf{W}_{down}) \mathbf{W}_{up} \quad (2)$$

247 Pfeiffer et al. (2020) have proposed a more efficient 290
 248 248
 249 249

250 3.2 Representation Editing

251 Previous PEFT methods fine-tune pre-trained mod- 291
 252 252
 253 253
 254 254
 255 255
 256 256
 257 257
 258 258
 259 259
 260 260

261 Turner et al. (2023) explicitly control the out- 292
 262 262
 263 263
 264 264
 265 265
 266 266
 267 267
 268 268
 269 269

270 Specifically, we first introduce a learnable scal- 293
 271 271
 272 272
 273 273
 274 274
 275 275
 276 276
 277 277
 278 278

$$279 \quad h_2 = l_{scaling} \odot h_1 + l_{bias} \quad (3)$$

280 ,where \odot denotes element-wise multiplication 294
 281 281
 282 282
 283 283

284 In addition, we initialize the scaling vector 295
 285 285
 286 286
 287 287
 288 288

4 Experiments 289

290 In this section, we conduct a series of experi- 290
 291 291
 292 292
 293 293
 294 294
 295 295
 296 296
 297 297
 298 298
 299 299
 300 300
 301 301
 302 302
 303 303
 304 304
 305 305
 306 306
 307 307
 308 308
 309 309

4.1 Baselines 310

311 To fully and fairly compare with other baselines, 311
 312 312
 313 313
 314 314
 315 315
 316 316

317 **Fine-Tuning (FT)** is a very common method for 317
 318 318
 319 319
 320 320
 321 321
 322 322
 323 323
 324 324
 325 325

326 **Bias-terms Fine-tuning (BitFit)** freezes most of 326
 327 327
 328 328
 329 329

330 **Adapter** adds the learnable lightweight module 330
 331 331
 332 332
 333 333
 334 334
 335 335
 336 336
 337 337
 338 338

Model & Method	# Params.	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	Avg.
FT (base)	125M	87.3	94.4	87.9	62.4	92.5	91.7	78.3	90.6	85.6
Adapter (base)	0.4M	87.0	93.3	88.4	60.9	92.5	90.5	76.5	90.5	85.0
LoRA (base)	0.3M	86.6	93.9	88.7	59.7	92.6	90.4	75.3	90.3	84.7
Adapter_FFN (base)	0.3M	87.1	93.0	88.8	58.5	92.0	90.2	77.7	90.4	84.7
BitFit (base)	0.1M	84.7	94.0	88.1	54.0	91.0	87.3	69.8	89.5	82.3
RED (base)	0.02M	83.9	93.9	89.2	61.0	90.7	87.2	78.0	90.4	84.3
FT (large)	355M	88.8	96.0	91.7	68.2	93.8	91.5	85.8	92.6	88.5
Adapter (large)	0.9M	90.1	95.2	90.5	65.4	94.6	91.4	85.3	91.5	88.0
LoRA (large)	0.8M	90.2	96.0	89.8	65.5	94.7	90.7	86.3	91.7	88.1
Adapter_FFN (large)	0.8M	90.3	96.1	90.5	64.4	94.3	91.3	84.8	90.2	87.7
RED (large)	0.05M	89.5	96.0	90.3	68.1	93.5	88.8	86.2	91.3	87.9

Table 1: RoBERTa base and RoBERTa large with RED and other adaptation methods on the GLUE benchmark. Higher is better for all metrics and fewer is better for the number of trainable parameters. The standard deviations of results from different methods are shown in Table 14 in Appendix C.

(2019).

Adapter_FFN is one kind of variant of Adapter proposed by Pfeiffer et al. (2020). Unlike the initial Adapter that requires inserting the learnable module between all sub-layers, Adapter_FFN only needs to apply an adapter after each FFN sub-layer.

AdapterDrop is another variant of Adapter proposed by Rücklé et al. (2020), which drops some adapter layers for greater efficiency.

Low-Rank Adaption(LoRA) performs low-rank decomposition on the incremental matrix ΔW and models the weight changes by multiplying two low-rank matrices. These two learnable matrices are concatenated in parallel next to the pre-trained model matrix, and they simultaneously process the input and add up the computation results as the output of this block, referred to Hu et al. (2021).

Prompt Tuning(PT) prefixes some continuous vectors at the embedding layer, which are learnable and are generally not in the model’s vocabulary, referred to Lester et al. (2021b).

Prefix tuning is a general version of prompt tuning, which prepends the continuous vectors at each hidden state, and these continuous vectors participate in the calculation of attention as key vectors and value vectors, referred to Li and Liang (2021b).

4.2 RoBERTa base/large

We take the pre-trained RoBERTa base (125M) and RoBERTa large (355M) from the HuggingFace

Transformers library (Wolf et al., 2019) and evaluate the performance of different efficient adaptation approaches on tasks from the GLUE benchmark, which is a widely recognized benchmark for natural language understanding. Moreover, we also replicate prior work according to their setup and conduct experiments under fair and reasonable configuration, see Appendix B.1 for more details of the hyperparameter used in our experiments.

Unlike previous works (Liu et al., 2019; Hu et al., 2021) that use the best model checkpoint on the MNLI dataset to initial model when dealing with MRPC, RTE, and STS-B to boost the performance, we consider a more general setting that trains the model from scratch.

The experimental performance of RED, as well as other adaption methods, is recorded in Table 1. Our results indicate that RED is comparable to other PEFT methods, which underlines the validity of directly editing representation as a feasible solution to adapt pre-trained models to downstream tasks.

RED has demonstrated strong competitiveness on training datasets with a data size of less than 100k, such as being able to match or even surpass other PEFT methods on SST-2, MRPC, CoLA, STS-B, and RTE. For datasets with data sizes greater than 100k, such as MNLI, QQP, and QNLI, the performance of BitFit and RED, which have the smallest number of trainable parameters, is slightly lower than other baselines. We think that larger-scale training datasets may require more trainable parameters to adapt.

Moreover, RED is highly parameter efficient. It can still maintain very good performance even

Model & Method	# Params.	BLEU	NIST	MET	ROUGE-L	CIDEr
FT (medium)	355M	65.95	8.52	45.95	69.13	2.35
FT^{top2} (medium)	25.2M	65.94	8.53	44.28	68.83	2.23
Adapter (medium)	0.9M	64.31	8.29	44.91	67.72	2.28
LoRA (medium)	0.8M	67.43	8.65	46.01	69.64	2.42
Adapter_FFN (medium)	0.8M	64.41	8.30	44.74	67.53	2.29
Prefix Tuning (medium)	0.8M	63.92	8.26	41.81	66.86	2.03
RED (medium)	0.05M	64.86	8.36	44.99	67.62	2.28
FT (large)	774M	65.56	8.50	45.40	68.38	2.27
Adapter (large)	1.8M	65.94	8.46	45.78	68.65	2.34
LoRA (large)	1.5M	68.24	8.76	46.23	69.92	2.42
Adapter_FFN (large)	1.5M	65.53	8.41	45.65	68.46	2.33
Prefix Tuning (large)	1.5M	65.50	8.45	43.97	67.32	2.23
RED (large)	0.09M	65.77	8.42	46.12	69.03	2.36

Table 2: GPT-2 medium and large with RED and other adaptation methods on the E2E NLG Challenge. Higher is better for all metrics and fewer is better for the number of trainable parameters. The standard deviations of results from different methods are shown in Table 15 in Appendix C.

with 7,200 times less trainable parameters than full parameter fine-tuning and 16 times less trainable parameters than LoRA, indicating that there is still room for improvement to further reduce the number of trainable parameters even if prior works adopt very sparse network structures to achieve this goal, which is consistent with the conclusions found by [Aghajanyan et al. \(2020\)](#) and [Kopiczko et al. \(2023\)](#).

4.3 GPT-2 medium/large

In addition to natural language understanding tasks, we also conduct experiments on natural language generation tasks. We take the pre-trained GPT-2 medium (355M) and GPT-2 large (774M) from the HuggingFace Transformers library and evaluate these methods on E2E NLG Challenge ([Novikova et al., 2017](#)). What’s more, we replicate prior works according to the setup of [Li and Liang \(2021a\)](#) and [Hu et al. \(2021\)](#), see Appendix B.2 for more details of the hyperparameter used in our experiments.

The experimental performance of RED, as well as other adaption methods, is recorded in Table 2. Our experimental results indicate that RED achieved comparable performance with other baselines in various metrics of the E2E NLG Challenge, proving that adapting downstream tasks through editing representation not only works on classification tasks but also performs well on generation tasks.

Similarly, RED achieves excellent performance with much less trainable parameters. In section

5.3, we set the rank of LoRA and Adapter to 1. At this point, the performance of RED is better than these PEFT methods, proving that RED is not only parameter efficient but also parameter effective.

4.4 T5

To further verify the universality of RED, we also select the encoder-decoder architecture model for experiments. Specifically, we take the pre-trained T5-base (220M) from the HuggingFace Transformers library and evaluate these methods on GLUE BenchMark. We reuse the results provided by [Asai et al. \(2022\)](#) and conduct experiments based on similar settings, see Appendix B.3 for more details of the hyperparameter used in our experiments.

The experimental performance of RED, as well as other adaption methods, is recorded in Table 3. Compared with other baselines, RED still achieved comparable results on T5 even with fewer parameter adjustments, indicating its universality and versatility for various model architectures.

RED demonstrated a significant performance advantage when compared to the PEFT method of Prompt Tuning, which has a similar number of trainable parameters. To be specific, RED surpasses Prompt Tuning by 11.9 points, indicating that allocating a small number of learnable parameters at each layer to edit representation is a more appropriate approach compared to allocating all learnable parameters on the embedding layer of the model.

Model & Method	# Params.	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	Avg.
FT (base)*	220M	86.8	94.6	90.2	61.8	93.0	91.6	71.9	89.7	84.9
Adapter (base)*	1.9M	86.5	93.8	85.3	64.0	93.2	90.2	71.9	90.7	84.5
AdapterDrop (base)*	1.1M	86.3	93.6	86.3	62.7	93.2	90.2	71.2	91.4	84.4
BitFit (base)*	0.3M	85.3	94.2	86.8	58.2	93.0	90.1	67.6	90.9	83.3
PT (base)*	0.08M	81.3	90.9	68.1	10.6	92.8	89.7	54.7	89.5	72.2
RED (base)	0.04M	85.9	93.0	91.7	61.1	91.2	89.2	72.7	88.2	84.1

Table 3: T5 base with RED and other adaptation methods on the GLUE benchmark. Higher is better for all metrics and fewer is better for the number of trainable parameters. * indicates numbers published in prior works.

4.5 Llama-2

Lastly, we scale up to Llama-2 with 7 billion parameters to verify the feasibility of applying RED on large-scale language models. Specifically, we selected Llama-2 (7B) as the base model and utilized full parameter fine-tuning, LoRA, and RED to fine-tune the model on UltraFeedback (Cui et al., 2023) respectively. As for evaluation, we assessed the performance of these different methods across three widely used benchmarks: Open LLM Leaderboard (Beeching et al., 2023), AlpacaEval (Li et al., 2023), and MT-Bench (Zheng et al., 2023). See Appendix B.4 for more details of the hyperparameter used in our experiments.

Method	# Params.	AlpacaEval (win %)
FT	6739M	80.93
LoRA	8.39M	81.48
RED	0.26M	81.69

Table 4: Win rates against reference response judged by GPT-4 on AlpacaEval. The higher the win rate, the more content generated is in line with human preferences, indicating better generation quality and fewer is better for the number of trainable parameters. See Figure 3 and Figure 4 in Appendix D for generated examples of RED compared with baselines in dialogue tasks.

Table 4 presents the win rates on AlpacaEval of responses generated by models trained with different methods, compared to the reference responses from text-davinci-003. RED achieved a higher win rate even though the number of trainable parameters was 25,700 times less than that of full parameter fine-tuning and 32 times less than LoRA, indicating that the method of directly editing representations to fine-tune the model is still applicable to large-scale language models and can generate the response that humans prefer.

Moreover, Figure 2 shows the performance score achieved by these adaption methods on 1-turn questions of MT-Bench. RED’s overall performance is

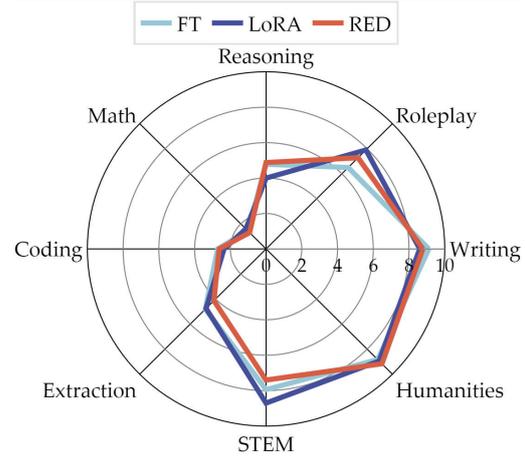


Figure 2: Scores of RED and other methods in MT-Bench. See more detailed results in Table 16 in Appendix D.

comparable to other baselines, and it has achieved the best results in evaluating the capabilities of Humanities and Reasoning. RED also achieved good results on six datasets of Open LLM Leaderboard, as shown in Table 18.

5 Ablation Study

In this section, we perform an ablation study to examine the impact of individual components of our method, editing representations in different positions and comparing the effectiveness of parameters between RED and other PEFT methods.

5.1 Contribution of different “edit vectors”

RED uses two different calculation types of “edit vectors”, scaling vector and bias vector, to edit representations. We remove the scaling vector and bias vector separately before editing the representation to explore the contribution of a single type of vector to this operation.

As shown in Table 6, when we remove any “edit vector” and then edit the representation, the performance on all datasets decreases to varying degrees,

Model & Method	# Params.	BLEU	NIST	MET	ROUGE-L	CIDEr
Adapter (rank 1)	0.25M	63.76	8.37	42.74	66.70	2.09
Adapter_FFNet (rank 1)	0.07M	62.99	8.09	40.88	66.39	1.98
LoRA (rank 1)	0.1M	64.51	8.38	44.78	67.35	2.28
RED	0.05M	64.86	8.36	44.99	67.62	2.28

Table 5: Comparison between RED and other rank 1 PEFT methods.

denoting that these two different vectors both have made contributions in the process of editing the representation. Compared to removing the scaling vector, removing the bias vector results in much more performance degradation, indicating that the bias vector plays a greater role in the process of editing representations.

Method	MRPC	CoLA	QQP
RED	90.3	68.1	88.7
-Scaling Vector	89.8	65.9	87.6
-Bias Vector	75.8	46.9	87.2

Table 6: Experimental results with different “edit vectors”. “-Scaling Vector” denotes that we remove the scaling vector and edit the representation by only bias vector; “-Bias Vector” denotes that we remove the bias vector and edit the representation by only scaling vector.

5.2 Position for editing representation

Houlsby et al. (2019) adds an Adapter module after both FFN and Attention sub-layers, Lin et al. (2020) only adds the Adapter module after the FFN sub-layer, and Hu et al. (2021) only adds the LoRA module in the Attention block, corresponding to different insertion positions of the PEFT component. To investigate the impact of operating different positions of the model on performance, we also designed experiments that only edited the representations after the FFN sub-layer, only edited the representations after the Attention sub-layer, and simultaneously edited the representations after the FFN and Attention sub-layer.

As shown in Table 7, editing only the representations after the FFN sub-layer yields slightly better performance compared to editing only the representations after the Attention sub-layer. Overall, there is not much change in performance compared to editing representations after both FFN and Attention sub-layer. Therefore, considering both performance and efficiency, a more favorable compromise entails restricting only editing the representations after the FFN sub-layer.

Position	MRPC	CoLA	QQP
FFN	90.3	68.1	88.7
Attn	88.9	66.1	88.6
FFN & Attn	89.2	68.4	89.4

Table 7: Experimental results on different positions.

5.3 Efficiency and effectiveness of parameters

When reproducing the experiments, we selected the rank of Adapter and LoRA according to the default settings of 8 or 16 in previous works, which may result in some parameter redundancy. Here, we set their rank to 1 and use GPT-2 medium as the base model to conduct experiments on the E2E NLG Challenge dataset, exploring the performance comparison between RED and these PEFT methods when their parameters are most efficient.

As shown in Table 5, compared to other variants with the smallest trainable parameters for each baseline, RED still has the smallest number of trainable parameters, indicating that RED is highly parameter efficient. Moreover, RED almost surpasses these baselines, which have slightly more trainable parameters, in all metrics, proving that RED is also parameter effective.

6 Conclusion

We explore fine-tuning the model from a new perspective of directly modifying the model representation, which is different from previous works that adjusted the model weights. We propose a new PEFT method **R**epresentation **E**ditting (RED), which fine-tunes the model by introducing two trainable “edit vectors” to edit representations. We have conducted extensive experiments on models of different architectures and scales on various types of NLP datasets, verifying that RED can still achieve comparable or even better performance than other baselines with much fewer trainable parameters, demonstrating that RED is not only parameter efficient but also parameter effective.

582 Limitations

583 We have demonstrated the effectiveness of the new
584 PEFT method of fine-tuning models by directly
585 editing the representations on various NLP tasks,
586 it would be intriguing to explore the application
587 of this method in other modalities, such as com-
588 puter vision and speech recognition. In addition,
589 articles related to representation engineering have
590 shown that only a very small number of examples
591 are needed to edit the representation to control the
592 model output. Therefore, we will also apply our
593 method to the few-shot scenarios to explore effec-
594 tive PEFT methods that are both parameter-efficient
595 and data-efficient in the future.

596 References

597 Armen Aghajanyan, Luke Zettlemoyer, and Sonal
598 Gupta. 2020. Intrinsic dimensionality explains the
599 effectiveness of language model fine-tuning. *arXiv*
600 *preprint arXiv:2012.13255*.

601 Akari Asai, Mohammadreza Salehi, Matthew E. Peters,
602 and Hannaneh Hajishirzi. 2022. *Attempt: Parameter-*
603 *efficient multi-task tuning via attentional mixtures of*
604 *soft prompts*. In *Conference on Empirical Methods*
605 *in Natural Language Processing*.

606 Satanjeev Banerjee and Alon Lavie. 2005. *Meteor: An*
607 *automatic metric for mt evaluation with improved*
608 *correlation with human judgments*. In *IEEvalu-*
609 *ation@ACL*.

610 Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro,
611 Danilo Giampiccolo, Bernardo Magnini, and Idan
612 Szpektor. 2006. The second pascal recognising tex-
613 tual entailment challenge.

614 Edward Beeching, Clémentine Fourrier, Nathan Habib,
615 Sheon Han, Nathan Lambert, Nazneen Rajani, Omar
616 Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023.
617 Open llm leaderboard. [https://huggingface.co/](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard)
618 [spaces/HuggingFaceH4/open_llm_leaderboard](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard).

619 Anja Belz and Ehud Reiter. 2006. *Comparing automatic*
620 *and human evaluation of nlg systems*. In *Conference*
621 *of the European Chapter of the Association for Com-*
622 *putational Linguistics*.

623 Elad Ben-Zaken, Shauli Ravfogel, and Yoav Gold-
624 berg. 2021. *Bitfit: Simple parameter-efficient*
625 *fine-tuning for transformer-based masked language-*
626 *models*. *ArXiv*, abs/2106.10199.

627 Tom Brown, Benjamin Mann, Nick Ryder, Melanie
628 Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind
629 Neelakantan, Pranav Shyam, Girish Sastry, Amanda
630 Askell, et al. 2020. Language models are few-shot
631 learners. *Advances in neural information processing*
632 *systems*, 33:1877–1901.

Daniel Matthew Cer, Mona T. Diab, Eneko Agirre, Iñigo
Lopez-Gazpio, and Lucia Specia. 2017. Semeval-
2017 task 1: Semantic textual similarity multilingual
and crosslingual focused evaluation. In *International*
Workshop on Semantic Evaluation. 633
634
635
636
637

Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao,
Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and
Maosong Sun. 2023. *Ultrafeedback: Boosting lan-*
guage models with high-quality feedback. *ArXiv*,
abs/2310.01377. 638
639
640
641
642

Dorottya Demszky, Kelvin Guu, and Percy Liang.
2018. Transforming question answering datasets
into natural language inference datasets. *ArXiv*,
abs/1809.02922. 643
644
645
646

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and
Luke Zettlemoyer. 2023. *Qlora: Efficient finetuning*
of quantized llms. *ArXiv*, abs/2305.14314. 647
648
649

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
Kristina Toutanova. 2018. Bert: Pre-training of deep
bidirectional transformers for language understand-
ing. *arXiv preprint arXiv:1810.04805*. 650
651
652
653

Ning Ding, Xingtai Lv, Qiaosen Wang, Yulin Chen,
Bowen Zhou, Zhiyuan Liu, and Maosong Sun. 2023.
Sparse low-rank adaptation of pre-trained language
models. *arXiv preprint arXiv:2311.11696*. 654
655
656
657

Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zong-
han Yang, Yusheng Su, Shengding Hu, Yulin Chen,
Chi-Min Chan, Weize Chen, et al. 2022. Delta tuning:
A comprehensive study of parameter efficient meth-
ods for pre-trained language models. *arXiv preprint*
arXiv:2203.06904. 658
659
660
661
662
663

William B. Dolan and Chris Brockett. 2005. *Automati-*
cally constructing a corpus of sentential paraphrases.
In *International Joint Conference on Natural Lan-*
guage Processing. 664
665
666
667

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman,
Sid Black, Anthony DiPofi, Charles Foster, Laurence
Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li,
Kyle McDonell, Niklas Muennighoff, Chris Ociepa,
Jason Phang, Laria Reynolds, Hailey Schoelkopf,
Aviya Skowron, Lintang Sutawika, Eric Tang, An-
ish Thite, Ben Wang, Kevin Wang, and Andy Zou.
2023. *A framework for few-shot language model*
evaluation. 668
669
670
671
672
673
674
675
676

Demi Guo, Alexander M Rush, and Yoon Kim. 2020.
Parameter-efficient transfer learning with diff prun-
ing. *arXiv preprint arXiv:2012.07463*. 677
678
679

Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-
Kirkpatrick, and Graham Neubig. 2021. *Towards a*
unified view of parameter-efficient transfer learning.
ArXiv, abs/2110.04366. 680
681
682
683

Dan Hendrycks, Collin Burns, Steven Basart, Andy
Zou, Mantas Mazeika, Dawn Xiaodong Song, and
Jacob Steinhardt. 2020. *Measuring massive multitask*
language understanding. *ArXiv*, abs/2009.03300. 684
685
686
687

688	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Xiaodong Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset . <i>ArXiv</i> , abs/2103.03874.		
689		Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models .	743
690			744
691			745
692			746
693	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 <i>International Conference on Machine Learning</i> , pages 2790–2799. PMLR.	Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	747
694			748
695			749
696		Stephanie C. Lin, Jacob Hilton, and Owain Evans. 2021. Truthfulqa: Measuring how models mimic human falsehoods . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	750
697			751
698			752
699	Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. <i>arXiv preprint arXiv:2106.09685</i> .		753
700		Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2020. Exploring versatile generative language model via parameter-efficient transfer learning . In <i>Findings</i> .	754
701			755
702			756
703		Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohhta, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning . <i>ArXiv</i> , abs/2205.05638.	757
704	Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021. Compacter: Efficient low-rank hypercomplex adapter layers. <i>Advances in Neural Information Processing Systems</i> , 34:1022–1035.		758
705			759
706			760
707			761
708	Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki Markus Asano. 2023. Vera: Vector-based random matrix adaptation . <i>ArXiv</i> , abs/2310.11454.	Wenhao Liu, Xiaohua Wang, Muling Wu, Tianlong Li, Changze Lv, Zixuan Ling, Jianhao Zhu, Cenyuan Zhang, Xiaoqing Zheng, and Xuanjing Huang. 2023. Aligning large language models with human preferences through representation engineering .	762
709			763
710			764
711	Jaejun Lee, Raphael Tang, and Jimmy J. Lin. 2019. What would elsa do? freezing layers during transformer fine-tuning . <i>ArXiv</i> , abs/1911.03090.		765
712			766
713		Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach . <i>ArXiv</i> , abs/1907.11692.	767
714	Brian Lester, Rami Al-Rfou, and Noah Constant. 2021a. The power of scale for parameter-efficient prompt tuning. <i>arXiv preprint arXiv:2104.08691</i> .		768
715			769
716			770
717	Brian Lester, Rami Al-Rfou, and Noah Constant. 2021b. The power of scale for parameter-efficient prompt tuning . In <i>Conference on Empirical Methods in Natural Language Processing</i> .		771
718		Yuning Mao, Lambert Mathias, Rui Hou, Amjad Almahairi, Hao Ma, Jiawei Han, Wen tau Yih, and Madian Khabsa. 2021. Unipelt: A unified framework for parameter-efficient language model tuning . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	772
719			773
720			774
721	Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario vSavsko, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clement Delangue, Théo Matussiere, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander M. Rush, and Thomas Wolf. 2021. Datasets: A community library for natural language processing . <i>ArXiv</i> , abs/2109.02846.		775
722			776
723		Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering . In <i>Conference on Empirical Methods in Natural Language Processing</i> .	777
724			778
725			779
726			780
727			781
728			782
729		Jekaterina Novikova, Ondrej Dusek, and Verena Rieser. 2017. The e2e dataset: New challenges for end-to-end generation . <i>ArXiv</i> , abs/1706.09254.	783
730			784
731			785
732		OpenAI. 2023. GPT-4 technical report . <i>CoRR</i> , abs/2303.08774.	786
733			787
734	Xiang Lisa Li and Percy Liang. 2021a. Prefix-tuning: Optimizing continuous prompts for generation. <i>arXiv preprint arXiv:2101.00190</i> .	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	788
735			789
736			790
737	Xiang Lisa Li and Percy Liang. 2021b. Prefix-tuning: Optimizing continuous prompts for generation . <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> , abs/2101.00190.		791
738			792
739		Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2020. Adapterfusion: Non-destructive task composition for transfer learning . <i>ArXiv</i> , abs/2005.00247.	793
740			794
741			795
742			

796	Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.	852
797		853
798		854
799	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners .	855
800		856
801		
802	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. <i>The Journal of Machine Learning Research</i> , 21(1):5485–5551.	
803		
804		
805		
806		
807		
808	Andreas Rücklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. 2020. Adapterdrop: On the efficiency of adapters in transformers . In <i>Conference on Empirical Methods in Natural Language Processing</i> .	857
809		858
810		859
811		860
812		861
813	Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavathula, and Yejin Choi. 2019. An adversarial winograd schema challenge at scale.	
814		
815		
816	Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Eliciting knowledge from language models using automatically generated prompts . <i>ArXiv</i> , abs/2010.15980.	862
817		863
818		864
819		865
820	Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, A. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank . In <i>Conference on Empirical Methods in Natural Language Processing</i> .	866
821		867
822		868
823		869
824		870
825		
826	Asa Cooper Stickland and Iain Murray. 2019. Bert and pals: Projected attention layers for efficient adaptation in multi-task learning. In <i>International Conference on Machine Learning</i> , pages 5986–5995. PMLR.	
827		
828		
829		
830		
831	Nishant Subramani, Nivedita Suresh, and Matthew E. Peters. 2022. Extracting latent steering vectors from pretrained language models . <i>ArXiv</i> , abs/2205.05124.	862
832		867
833		868
834	Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashii Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin	869
835		890
836		891
837		892
838		893
839		894
840		
841		
842		
843		
844		
845		
846		
847		
848		
849		
850		
851		
	Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models . <i>ArXiv</i> , abs/2307.09288.	852
		853
		854
		855
		856
	Alexander Matt Turner, Lisa Thiergart, David S. Udell, Gavin Leech, Ulisse Mini, and Monte Stuart MacDiarmid. 2023. Activation addition: Steering language models without optimization . <i>ArXiv</i> , abs/2308.10248.	857
		858
		859
		860
		861
	Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need . In <i>Neural Information Processing Systems</i> .	862
		863
		864
		865
	Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2014. Cider: Consensus-based image description evaluation . <i>2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 4566–4575.	866
		867
		868
		869
		870
	Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding . In <i>BlackboxNLP@EMNLP</i> .	871
		872
		873
		874
		875
	Zhen Wang, Rameswar Panda, Leonid Karlinsky, Rogério Schmidt Feris, Huan Sun, and Yoon Kim. 2023. Multitask prompt tuning enables parameter-efficient transfer learning . <i>ArXiv</i> , abs/2303.02861.	876
		877
		878
		879
	Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2018. Neural network acceptability judgments . <i>Transactions of the Association for Computational Linguistics</i> , 7:625–641.	880
		881
		882
		883
	Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. In <i>North American Chapter of the Association for Computational Linguistics</i> .	884
		885
		886
		887
		888
	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing . <i>ArXiv</i> , abs/1910.03771.	889
		890
		891
		892
		893
		894
	Muling Wu, Wenhao Liu, Jianhan Xu, Changze Lv, Zixuan Ling, Tianlong Li, Longtao Huang, Xiaoqing Zheng, and Xuanjing Huang. 2023. Parameter efficient multi-task fine-tuning by learning to transfer token-wise prompts . In <i>Conference on Empirical Methods in Natural Language Processing</i> .	895
		896
		897
		898
		899
		900
	Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models . <i>arXiv preprint arXiv:2106.10199</i> .	901
		902
		903
		904

905 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali
906 Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a ma-](#)
907 [chine really finish your sentence?](#) In *Annual Meeting*
908 *of the Association for Computational Linguistics*.

909 Qingru Zhang, Minshuo Chen, Alexander Bukharin,
910 Pengcheng He, Yu Cheng, Weizhu Chen, and
911 Tuo Zhao. 2023a. Adaptive budget allocation
912 for parameter-efficient fine-tuning. *arXiv preprint*
913 *arXiv:2303.10512*.

914 Qingru Zhang, Minshuo Chen, Alexander W. Bukharin,
915 Pengcheng He, Yu Cheng, Weizhu Chen, and
916 Tuo Zhao. 2023b. [Adaptive budget alloca-](#)
917 [tion for parameter-efficient fine-tuning.](#) *ArXiv*,
918 *abs/2303.10512*.

919 Mengjie Zhao, Tao Lin, Martin Jaggi, and Hinrich
920 Schütze. 2020. [Masking as an efficient alternative to](#)
921 [finetuning for pretrained language models.](#) In *Con-*
922 *ference on Empirical Methods in Natural Language*
923 *Processing*.

924 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan
925 Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin,
926 Zhuohan Li, Dacheng Li, Eric P. Xing, Haotong
927 Zhang, Joseph Gonzalez, and Ion Stoica. 2023. [Judg-](#)
928 [ing llm-as-a-judge with mt-bench and chatbot arena.](#)
929 *ArXiv*, *abs/2306.05685*.

930 Andy Zou, Long Phan, Sarah Chen, James Campbell,
931 Phillip Guo, Richard Ren, Alexander Pan, Xuwang
932 Yin, Mantas Mazeika, Ann-Kathrin Dombrowski,
933 Shashwat Goel, Nathaniel Li, Michael J. Byun,
934 Zifan Wang, Alex Mallen, Steven Basart, Sanmi
935 Koyejo, Dawn Song, Matt Fredrikson, Zico Kolter,
936 and Dan Hendrycks. 2023. [Representation engineer-](#)
937 [ing: A top-down approach to ai transparency.](#) *ArXiv*,
938 *abs/2310.01405*.

A Dataset And Evaluation Details

A.1 GLUE Benchmark

The GLUE benchmark, consisting of CoLA (Warstadt et al., 2018), SST-2 (Socher et al., 2013), MRPC (Dolan and Brockett, 2005), QQP (Wang et al., 2018), STS-B (Cer et al., 2017), MNLI (Williams et al., 2017), QNLI (Demszky et al., 2018) and RTE (Bar-Haim et al., 2006), is used for natural language understanding. We source each dataset from Huggingface Datasets (Lhoest et al., 2021) and utilize the full dataset for our experiments.

Following (Ding et al., 2023) and (Hu et al., 2021), we evaluate models on the validation dataset. But unlike Hu et al. (2021) which just uses the training dataset for training and the validation dataset for testing and selects the best result for each run, we have considered a more reasonable setting by dividing the validation set into validation set and test set. After each epoch training is completed, we will verify it on the validation set and record the verification results, after training all epochs, we select the model with the best performance on the validation set and test it on the test set. For datasets with a large validation set, we select 1000 samples as the validation set, and then use the remaining samples as the test set, and for datasets with a small validation set, we select half of the samples as the validation set, and then use the remaining samples as the test set, the details, and the evaluation metric are reported in Table 8.

For all experiments on RoBERTa, we run 5 times using different random seeds and report the average results in order to ensure statistical significance. To be specific, we use 42, 43, 44, 45, 46 these 5 random seeds².

A.2 E2E NLG Challenge

E2E NLG Challenge was first introduced in Novikova et al. (2017) as a dataset for training end-to-end, data-driven natural language generation systems and is commonly used for data-to-text evaluation.

We source each dataset from Huggingface Datasets and utilize the full dataset for our experiments. Specifically, this dataset contains 42.1k training samples, 4.67k validation samples, and 4.69k testing samples. Following previous works,

²When conducting experiments on the RTE dataset, some random seeds corresponded to abnormal experimental results, so several random seeds were replaced.

Dataset	#Train	#Valid	#Test	Metric
CoLA	8.5K	522	521	Mcc
SST-2	67k	436	436	Acc
MRPC	3.7K	204	204	Acc
QQP	364K	1K	39K	Acc
STS-B	5.7k	750	750	Corr
MNLI	393k	1K	8K	Acc
QNLI	105K	1K	4.5K	Acc
RTE	2.5k	139	138	Acc

Table 8: The size and evaluation metric of the split version datasets in GLUE benchmark. "Mcc", "Acc" and "Corr" represent Matthews correlation coefficient, accuracy, and Pearson correlation coefficient respectively.

we use the official evaluation script, which reports BLEU (Papineni et al., 2002), NIST (Belz and Reiter, 2006), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004) and CIDEr (Vedantam et al., 2014).

For all experiments on GPT-2, we run 3 times using different random seeds and report the average results in order to ensure statistical significance. To be specific, we use 42, 43, and 44 these 3 random seeds.

A.3 UltraFeedback

UltraFeedback (Cui et al., 2023) consists of 64k prompts, each of which has four LLM responses that are rated by GPT-4 according to criteria like instruction-following, honesty, and helpfulness. We construct our training dataset from UltraFeedback by selecting the highest mean score as the "chosen" response.

A.4 Open LLM Leaderboard

Open LLM Leaderboard comprises six benchmarks that cover science questions, commonsense inference, multitask accuracy, math reasoning, and truthfulness in generating answers. Specifically, it consists of ARC (Mihaylov et al., 2018), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2019), MMLU (Hendrycks et al., 2020), TruthfulQA (Lin et al., 2021), and GSM8K (Hendrycks et al., 2021). We utilized the Eleuther AI Language Model Evaluation Harness library (Gao et al., 2023) to assess language models trained using different methods. Table 17 provides a detailed description of the leaderboard evaluation configuration and the experimental settings adopted in this study.

Method & Model	Dataset	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B
	Optimizer	AdamW							
	LR Schedule	Linear							
FT	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning Rate	$1e-05$	$1e-05$	$1e-05$	$1e-05$	$1e-05$	$1e-05$	$1e-05$	$1e-05$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Weight Decay	$1e-04$							
	Max Seq. Len.	256							
LoRA	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning Rate	$5e-04$	$5e-04$	$4e-04$	$4e-04$	$4e-04$	$5e-04$	$5e-04$	$4e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	LoRA Config.	$r_q = r_v = 8$							
	LoRA α .	8							
Adapter	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning Rate	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Rank.	8							
	Max Seq. Len.	256							
Adapter_FFNN	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning Rate	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Rank.	16							
	Max Seq. Len.	256							
BitFit	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning Rate	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$	$1e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Max Seq. Len.	256							
RED	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	20	40	40	40	20	20	40	40
	Learning rate	$1e-03$	$1e-03$	$5e-03$	$5e-03$	$3e-03$	$1e-03$	$8e-03$	$3e-03$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Max Seq. Len.	256							

Table 9: Hyperparameters for RoBERTa based on GLUE.

Method & Model	Dataset	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B
	Optimizer	AdamW							
	LR Schedule	Linear							
FT	Batch Size	16	32	32	32	16	16	16	32
	# Epochs	10	10	20	20	10	10	20	10
	Learning rate	$2e-05$	$2e-05$	$2e-05$	$2e-05$	$2e-05$	$1e-05$	$1e-05$	$2e-05$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Weight Decay	$1e-01$							
	Max Seq. Len.	256							
LoRA	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	10	10	20	20	10	10	20	10
	Learning rate	$3e-04$	$4e-04$	$3e-04$	$2e-04$	$2e-04$	$3e-04$	$4e-04$	$2e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	LoRA Config.	$r_q = r_v = 8$							
	LoRA α .	16							
Max Seq. Len.	256								
Adapter	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	10	10	20	20	10	10	20	10
	Learning rate	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Rank.	8							
	Max Seq. Len.	256							
Adapter_FFNet	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	10	10	20	20	10	10	20	10
	Learning rate	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$	$3e-04$
	Warmup Ratio	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$	$6e-02$
	Rank.	16							
	Max Seq. Len.	256							
RED	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	10	10	20	20	10	10	20	10
	Learning rate	$1e-03$	$1e-03$	$2e-03$	$1e-03$	$1e-03$	$1e-03$	$5e-03$	$5e-03$
	Weight Decay	0.0	0.0	0.0	0.0	0.0	0.0	$1e-04$	0.0
	Warmup Ratio	$6e-02$	$6e-02$	0	$6e-02$	$6e-02$	$6e-02$	$1e-02$	$6e-02$
	Max Seq. Len.	256							

Table 10: Hyperparameters for RoBERTa large on GLUE.

Dataset	E2E NLG Challenge							
	Training							
	FT	FT_top2	Adapter	Apapter_FFNet	LoRA	Prefix Tuning	RED_M	RED_L
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Weight Decay	0.0	0.0	0.0	0.0	$1e-02$	0.0	$1e-04$	0.0
# Epoch	5	5	5	5	5	5	5	10
Learning Rate Schedule	Linear	Linear	Linear	Linear	Linear	Linear	Linear	Linear
Label Smooth	0.0	0.0	0.0	0.0	0.1	0.0	0.0	0.0
Learning Rate	$5e-05$	$5e-05$	$8e-05$	$8e-05$	$2e-04$	$8e-05$	$6e-02$	$5e-05$
Rank or Prefix Length	-	-	8	16	8	16	-	-
Lora α	-	-	-	-	32	-	-	-
Adaption	-	-	-	-	$r_q = r_v = r_k = 8$	-	-	-
Warmup Steps	500	500	500	500	500	500	500	500
Batch Size	10	10	10	10	10	10	10	10
	Inference							
Beam Size	10							
Length Penalty	0.9							
no repeat ngram size	4							

Table 11: Hyperparameters for GPT-2 on E2E NLG Challenge. RED_M represents hyperparameters applied on GPT-2 medium and RED_L represents hyperparameters applied on GPT-2 large.

Method & Model	Dataset	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B
	Optimizer	AdamW							
	LR Schedule	Linear							
	Batch Size	32	32	32	32	32	32	32	32
	# Epochs	10	10	20	60	10	10	30	50
RED	Learning rate	$5e-02$	$3e-02$	$1e-01$	$4e-02$	$2e-02$	$5e-03$	$7e-02$	$4e-02$
	Warmup Ratio	$1e-02$	$6e-02$	$1e-02$	$1e-02$	$0e-00$	$1e-02$	$1e-02$	$6e-02$
	Max Seq. Len.	256							

Table 12: Hyperparameters for T5 base on GLUE.

Method	Hyperparameter	Value
	Batch Size	128
	Micro Batch Size	2
	Optimizer	Adamw
	LR Scheduler Type	Cosine
	Warmup Ratio	0.1
	Max Seq. Len.	768
FT	Learning Rate	$2e-5$
	# Epochs	3
	Learning Rate	$3e-4$
	# Epochs	3
LoRA	Batch Size	128
	LoRA α	16
	LoRA Dropout	0.1
	LoRA Rank	16
	Target Modules	[q_proj, v_proj]
RED	Learning Rate	$1e-3$
	# Epochs	10

Table 13: Hyperparameters used for Llama-2.

Model & Method	# Params.	MNLI	SST-2	MRPC	CoLA	QNLI	QQP	RTE	STS-B	Avg.
FT (base)	125M	87.3 (0.34)	94.4 (0.96)	87.9 (0.91)	62.4 (3.29)	92.5 (0.22)	91.7 (0.19)	78.3 (3.20)	90.6 (0.59)	85.6
Adapter (base)	0.4M	87.0 (0.28)	93.3 (0.40)	88.4 (1.54)	60.9 (3.09)	92.5 (0.02)	90.5 (0.08)	76.6 (2.26)	90.5 (0.35)	85.0
Adapter_FFN (base)	0.3M	87.1 (0.10)	93.0 (0.50)	88.8 (1.38)	58.5 (1.69)	92.1 (0.28)	90.2 (0.07)	77.7 (1.93)	90.4 (0.31)	84.7
LoRA (base)	0.3M	86.6 (0.23)	93.9 (0.49)	88.7 (0.76)	59.7 (4.36)	92.6 (0.10)	90.4 (0.08)	75.3 (2.79)	90.3 (0.54)	84.7
BitFit (base)	0.1M	84.7 (0.08)	94.0 (0.87)	88.1 (1.57)	54.0 (3.07)	91.0 (0.05)	87.3 (0.02)	69.8 (1.51)	89.5 (0.35)	82.3
RED (base)	0.02M	83.9 (0.14)	93.9 (0.31)	89.2 (0.98)	61.0 (2.96)	90.7 (0.35)	87.2 (0.17)	78.0 (2.06)	90.4 (0.32)	84.7
FT (large)	355M	88.8 (0.45)	96.0 (0.66)	91.7 (1.73)	68.2 (2.62)	93.8 (0.33)	91.5 (1.28)	85.8 (1.40)	92.6 (0.16)	88.5
LoRA (large)	0.8M	90.2 (0.25)	96.0 (0.85)	89.8 (2.09)	65.5 (2.02)	94.7 (0.21)	90.7 (0.91)	86.3 (2.41)	91.7 (0.44)	88.1
Adapter (large)	0.9M	90.1 (0.12)	95.2 (0.48)	90.5 (0.59)	65.4 (2.24)	94.6 (0.17)	91.4 (0.13)	85.3 (1.34)	91.5 (0.33)	88.0
Adapter_FFN (large)	0.8M	90.3 (0.15)	96.1 (0.75)	90.5 (1.26)	64.4 (1.56)	94.3 (0.39)	91.3 (0.24)	84.8 (2.01)	90.2 (0.24)	87.7
RED (large)	0.05M	89.5 (0.38)	96.0 (0.48)	90.3 (1.40)	68.1 (1.69)	93.5 (0.33)	88.8 (0.11)	86.2 (1.40)	91.3 (0.21)	87.9

Table 14: Test results of RED and other adaption methods on the GLUE benchmark. The standard deviation is provided in parentheses.

Model & Method	# Params.	BLEU	NIST	MET	ROUGE-L	CIDEr
FT (medium)	355M	65.95 (0.26)	8.52 (0.03)	45.95 (0.07)	69.13 (0.30)	2.35 (0.01)
FT^{top2} (medium)	25.2M	65.95 (0.33)	8.52 (0.03)	45.95 (0.09)	69.13 (0.17)	2.35 (0.02)
Adapter (medium)	0.9M	64.31 (0.17)	8.29 (0.01)	44.91 (0.29)	67.72 (0.26)	2.28 (0.01)
Adapter_FFN (medium)	0.8M	64.41 (0.17)	8.30 (0.02)	44.74 (0.11)	67.53 (0.02)	2.29 (0.01)
LoRA (medium)	0.8M	67.43 (0.39)	8.65 (0.05)	46.01 (0.07)	69.64 (0.14)	2.42 (0.01)
Prefix Tuning (medium)	0.8M	63.92 (0.27)	8.26 (0.11)	41.81 (0.62)	66.86 (0.22)	2.03 (0.05)
RED (medium)	0.05M	64.86 (0.40)	8.36 (0.03)	44.99 (0.02)	67.62 (0.22)	2.28 (0.01)
FT (large)	774M	65.56 (0.47)	8.50 (0.05)	45.40 (0.29)	68.38 (0.23)	2.27 (0.02)
Adapter (large)	1.8M	65.94 (0.35)	8.46 (0.05)	45.78 (0.11)	68.65 (0.35)	2.34 (0.01)
Adapter_FFN (large)	1.5M	65.53 (0.61)	8.41 (0.07)	45.65 (0.12)	68.46 (0.16)	2.33 (0.01)
LoRA (large)	1.5M	68.24 (0.28)	8.76 (0.04)	46.23 (0.04)	69.92 (0.16)	2.42 (0.01)
Prefix Tuning (large)	1.5M	65.50 (0.63)	8.45 (0.05)	43.97 (0.21)	67.32 (0.38)	2.23 (0.02)
RED (large)	0.09M	65.77 (0.48)	8.42 (0.06)	46.12 (0.10)	69.03 (0.09)	2.36 (0.02)

Table 15: Test results of RED and other adaption methods on the E2E NLG Challenge dataset. The standard deviation is provided in parentheses.

Method	Trainable Params.	Writing	Roleplay	Reasoning	Math	Coding	Extraction	Stem	Humanities	Average
Turn-1										
FT	6739M	9.111	6.500	4.778	1.444	2.700	4.800	7.944	8.833	5.688
LoRA	8.39M	8.600	7.900	4.000	1.600	2.350	4.750	8.700	8.950	5.856
RED	0.26M	8.778	7.278	4.875	1.300	2.625	4.125	7.400	9.167	5.732
Turn-2										
FT	6739M	1.667	5.938	2.222	1.700	2.000	2.111	3.200	5.300	3.021
LoRA	8.39M	1.900	5.800	2.100	1.600	2.550	1.222	3.100	5.500	2.994
RED	0.26M	1.375	5.500	2.444	1.444	2.125	1.750	3.000	5.875	2.946
Final										
FT	6739M	5.389	6.200	3.500	1.579	2.412	3.526	5.447	6.974	4.354
LoRA	8.39M	5.250	6.850	3.050	1.600	2.450	3.079	5.900	7.225	4.434
RED	0.26M	5.294	6.389	3.588	1.368	2.375	2.938	5.750	7.618	4.401

Table 16: Results of MT-Bench.

Datasets	Arc	TruthfulQA	Winogrande	GSM8k	HellaSwag	MMLU
# few-shot	25	0	5	5	10	5
Metric	acc_norm	mc2	acc	acc	acc_norm	acc

Table 17: For each dataset used in the evaluation on the Open LLM Leaderboard, we detail the quantity of few-shot samples utilized and the specific metric employed for evaluation.

Method	# Params.	Arc	TruthfulQA	Winogrande	GSM8k	HellaSwag	MMLU	Average
FT	6739M	73.34	47.07	74.03	16.22	79.95	46.55	56.19
LoRA	8.39M	72.07	44.84	75.85	15.01	78.60	44.42	55.13
RED	0.26M	72.04	47.59	72.85	11.90	77.86	42.27	54.09

Table 18: Results of different methods on six benchmarks of Open LLM Leaderboard.

Prompt Can you help me figure out how to sew a button on a shirt?	
FT	<ol style="list-style-type: none"> 1. Gather the materials needed for sewing a button, which include thread, needle, and button. 2. Thread the needle with the desired color of thread. 3. Place the button through the hole in the shirt where it is missing or needs replacing. 4. Holding the button in place, thread the needle through the fabric from the backside of the shirt. 5. Pull the thread tightly to secure the button in place. 6. Cut off any excess thread. 7. Sew the button onto the shirt by hand or using a sewing machine. 8. Iron the shirt to smooth out any creases caused during the sewing process. 9. Enjoy your newly repaired shirt!
LoRA	<ol style="list-style-type: none"> 1. Gather the necessary materials: You will need a shirt, thread, a needle, and a button. If your shirt has a buttonhole, you can use that; otherwise, you'll have to create one using a buttonhole template or a pattern. 2. Mark the location of the buttonhole: Use a marker or chalk to mark the location where you want to place the buttonhole. This should be centered on the front of the shirt, just below the collar. 3. Create the buttonhole: If you don't have a buttonhole already, you'll need to make one. To do this, fold the fabric over so that it forms a triangle shape. Then, use a ruler and a pencil to draw a line along the edge of the triangle. Cut along this line with scissors. 4. Sew the buttonhole: Thread the needle with the thread and insert the needle through the fabric at the top of the buttonhole. Pull the thread tightly to create a small hole in the fabric. Continue sewing around the perimeter of the buttonhole, leaving about an inch of thread unsewn at the bottom. 5. Insert the button: Place the button on the shirt, right-side up, with the hole facing down. Line up the buttonhole with the hole in the button and push the button through the buttonhole. 6. Tie off the thread: Once the button is securely attached, tie off the thread by knotting it near the beginning of the buttonhole stitching. Trim any excess thread. 7. Finish the buttonhole: Fold the fabric over the buttonhole and press it flat with an iron. This will help to secure the button and prevent it from coming loose. 8. Repeat for additional buttons: If you need to add more buttons, repeat steps 4-7 for each new button. 9. Test fit: Try on the shirt to ensure that the buttons are evenly spaced and that they look good. Adjust as needed. 10. Enjoy your newly buttoned shirt!
Red	<ol style="list-style-type: none"> 1. Gather the necessary materials: You will need a shirt, thread, a needle, and a button. If your shirt has buttons that are already missing or damaged, you can purchase replacement buttons at most craft stores. 2. Prepare the shirt: Lay the shirt flat with the back facing upwards. Make sure the area where you want to sew the button is clean and free of any debris or wrinkles. 3. Thread the needle: Insert one end of the thread into the eye of the needle and pull it through the other side. Tie a knot in the end of the thread to secure it. 4. Place the button on the shirt: Position the button on the desired location on the shirt, making sure it's centered and aligned properly. 5. Sew the button: Holding the shirt steady with one hand, use the other hand to guide the needle through the fabric from the backside of the shirt. Inject the thread into the fabric by pushing down gently with your thumb while pulling the needle through. Continue sewing around the perimeter of the button, making sure to keep the stitches even and consistent. 6. Trim the excess thread: Once you have finished sewing the button, cut off any excess thread that extends beyond the edge of the button. 7. Test the button: Try fastening the button onto the shirt to make sure it's securely attached. If needed, adjust the position of the button or re-sew if it's not properly attached. 8. Repeat for additional buttons: If there are multiple buttons on the shirt that need to be sewn, repeat steps 1-7 for each button. 9. Clean up: After completing all the buttons, carefully remove any loose threads or lint from the shirt. <p>Remember to always use caution when handling sharp objects like needles and thread. Always wear protective gloves and follow proper safety precautions to avoid injury.</p>

Figure 3: Response from RED provides a comprehensive, step-by-step guide that is detailed and methodical, ensuring that even beginners can follow along easily. It covers preparation, threading, placement, sewing technique, and cleanup, offering safety tips and encouraging testing for security. Conversely, the response from LoRA focuses incorrectly on creating buttonholes rather than sewing buttons, which might confuse readers seeking to attach buttons. Response from FT offers a simplified overview but lacks the depth and clarity of RED, making it less instructive for those unfamiliar with sewing. RED's detailed approach ensures accuracy and success in sewing a button.

Prompt	<p>Hi, my cat just chewed on a small weed outside that looked like hairy bittercress. Could you check if this is toxic for cats? Please also look for plants from the same family, as I'm not 100% sure I identified it correctly.</p>
FT	<p> Hello! I'm here to help answer your question about whether or not the plant your cat chewed on is toxic for them. However, before I can provide an accurate response, I need more information about the plant in question. Could you please describe the plant in detail, including its appearance, size, and any unique characteristics it may have? Additionally, could you tell me what part of the plant your cat ingested (e.g., leaves, stems, flowers)? This will help me identify the correct plant and determine if it is safe or harmful for cats.</p> <p>As a responsible AI assistant, I want to ensure that I provide accurate and helpful information while prioritizing the safety and well-being of both you and your furry friend. So, please take a moment to share any additional details you can about the plant your cat has chewed on, and I'll do my best to assist you.</p>
LoRA	<p> Hello! I'm here to help answer your question about whether or not the plant your cat chewed on is toxic for them. However, before I can provide an accurate response, I need some more information. Can you please tell me what kind of plant it was and where it was found? This will help me identify the species and determine if it's safe for cats to consume. Additionally, I would appreciate it if you could provide any other details you have about the plant, such as its appearance, smell, and location. Thank you for your patience and understanding!</p>
Red	<p> Hello! I'm here to help you with your question about whether or not your cat has eaten something potentially harmful. As an AI language model, I can provide information and resources based on the best available evidence. However, it's important to note that I am not a veterinarian and cannot diagnose or treat any medical conditions in animals. If you have concerns about your pet's health, please consult a licensed veterinarian.</p> <p>Now, let's focus on your specific question. Based on what you described, it sounds like your cat may have eaten some type of plant material. It's always a good idea to keep an eye on your pet after they eat something new, especially if it's a plant. Some common signs of poisoning in cats include vomiting, diarrhea, drooling, lethargy, weakness, seizures, and difficulty breathing. If you notice any of these symptoms, it's essential to seek immediate veterinary care.</p> <p>As for the plant itself, it's difficult to identify without more details. Bittercress (<i>Ranunculus</i> spp.) is a genus of flowering plants that includes several species, including buttercup, crowfoot, and spearwort. These plants are generally considered safe for cats, but there are other members of the Ranunculaceae family that could be toxic, such as aconite (<i>Aconitum</i> spp.), monkshood (<i>Aconitella</i> spp.), and wolfsbane (<i>Aconitine</i>).</p> <p>To be on the safe side, I recommend taking your cat to the vet for a thorough examination and possible testing if necessary. The vet will be able to determine if the plant your cat ate was indeed toxic and provide appropriate treatment if needed. They may also advise you on how to prevent future exposure to potentially harmful plants.</p> <p>In conclusion, while most plants are safe for cats to consume, it's always better to err on the side of caution and seek professional advice when in doubt. Please take your cat to the veterinarian for a proper evaluation and follow their instructions accordingly.</p>

Figure 4: Response from RED offers a comprehensive and proactive approach, addressing immediate concerns, potential symptoms to watch for, and the importance of veterinary consultation. It provides a general safety assessment of the Ranunculaceae family, including possible toxic members, emphasizing the need for professional evaluation. This response balances informative content with practical advice, ensuring pet owners are prepared to act in their pet's best interest, even without specific plant identification. In contrast, responses from FT and LoRA focus more on gathering additional information before offering advice, which could delay essential care in a potentially urgent situation.