

Context-aware Prompt Tuning: Advancing In-Context Learning with Adversarial Methods

Anonymous ACL submission

Abstract

Large Language Models (LLMs) can perform few-shot learning using In-Context Learning (ICL) or optimization-based methods. ICL is more effective in low-data regimes, while optimization-based methods excel with larger datasets. This contrast raises a key question: why optimization-based methods face challenges in low-data regimes, and how can these methods be effectively integrated with ICL to enhance few-shot learning? In this work, we identify overfitting as the primary limitation of optimization-based methods in few-shot settings and introduce Context-aware Prompt Tuning (CPT), a method that combines the strengths of ICL, Prompt Tuning (PT), and adversarial techniques. CPT initializes the context with training examples, similar to ICL, and then applies an optimization process inspired by PT and adversarial techniques. Through iterative adaptation, CPT effectively balances flexibility and stability, allowing it to derive deeper insights from limited data while preserving the integrity of input samples. Our method achieves superior accuracy across multiple classification tasks and LLM architectures, consistently outperforming existing baselines and effectively mitigating overfitting challenges in few-shot scenarios.

1 Introduction

Adapting Large Language Models (LLMs) to new tasks in few-shot learning scenarios can be achieved through either fine-tuning or In-Context Learning (ICL) (Brown et al., 2020). Parameter-efficient fine-tuning methods, such as Low-Rank Adaptation (LoRA) (Hu et al., 2021), which optimizes a subset of the model’s parameters, and Prompt Tuning (PT) (Lester et al., 2021), which optimizes a small set of learnable tokens prepended to the input, aim to achieve task-specific performance with minimal computational overhead. In contrast, ICL eliminates the need for parameter updates by




Figure 1: **Overfitting Across Few-Shot Methods** Train-test loss gap across methods and training set sizes using the GPT-J model on the DBpedia dataset. For each model, there are two loss graphs: one for train loss (dotted line) and one for test loss (solid line). CPT performs better in mitigating overfitting compared to optimization-based methods. Despite a relatively higher training loss, CPT achieves the lowest test loss.

incorporating training examples directly into the input context, offering a training-free alternative that leverages the model’s pre-trained knowledge without modifying its underlying parameters. Despite their effectiveness, determining the optimal approach for varying dataset sizes remains an ongoing challenge.

In few-shot scenarios with limited data, ICL has shown greater effectiveness; however, as the dataset size increases, optimization-based methods like LoRA and PT become preferable. This trend has been observed in prior studies (Mosbach et al., 2023; Min et al., 2022) and is further supported by our experimental results (fig. 3), which demonstrate that while ICL excels in low-data regimes, its advantage diminishes as more data becomes available. Although promising, the limitations of optimization-based methods in low-data scenarios require further exploration.

In this work, we identify overfitting as the primary factor limiting the effectiveness of optimization-based methods in few-shot learning scenarios, as demonstrated in fig. 1. To address this challenge, we propose Context-Aware Prompt Tuning (CPT), a novel approach that integrates concepts from ICL, Prompt Tuning (PT), and adversarial attacks (Blau et al., 2022, 2023; Carlini and Wagner, 2017; Athalye et al., 2018; Madry




Figure 2: **Comparison of Few-Shot Methods** We highlight the key differences between *CPT* and the baseline methods, focusing on *ICL*, *PT*, and *IPT*. Each method includes two token types: prefix tokens (blue background) and loss tokens (orange background). The tokens are categorized into *Learnable Tokens* (pink) and *Sample Tokens* (brown), which remain fixed during training. A red line beneath the tokens indicates those used for loss calculation, while a green line marks those updated during training. *CPT* introduces *Learnable Sample* tokens, shown in a brown-pink color, initialized with training samples and progressively refined during optimization.

et al., 2017; Goyal et al., 2020). As illustrated in fig. 2, CPT utilizes training examples in two ways: first, to construct the context as in ICL, and second, to optimize the context token embeddings using methodologies inspired by PT and adversarial attacks. To combat overfitting, CPT refines the context tokens while preserving their structure, integrates context labels into the loss function as a form of regularization, and applies projected gradient descent to maintain proximity to their original values. Additionally, CPT employs a loss weighting mechanism that leverages recency bias—a phenomenon where models prioritize later examples in the context (Zhao et al., 2021), thereby guiding the model to prioritize the most relevant examples during optimization. These strategies collectively enable CPT to strike a balance between optimization flexibility and robustness, effectively addressing overfitting challenges in few-shot learning scenarios.

We rigorously evaluate CPT across multiple classification tasks and model architectures, conducting extensive ablation studies to validate each design choice. To ensure robustness, we employ diverse templates and seeds—an essential considering ICL’s sensitivity to the selection and formatting of training examples, as highlighted by (Sun et al., 2023; Zhao et al., 2021). Our results demon-

strate that CPT consistently achieves superior performance compared to existing baselines across diverse scenarios.

To summarize, our key contributions are as follows:

- We identify overfitting as the primary limitation of optimization-based methods in few-shot learning scenarios and empirically demonstrate its impact on performance degradation.
- We propose Context-Aware Prompt Tuning (CPT), a novel few-shot learning method that enhances ICL with optimization-based techniques. CPT employs targeted strategies to effectively mitigate overfitting.
- We achieve state-of-the-art results across multiple classification datasets and perform extensive ablation studies to validate each component of our proposed method.

2 Related Work

Fine-Tuning Fine-tuning is a popular and effective method for adjusting LLMs to specific tasks. Standard fine-tuning (Radford et al., 2019; Brown et al., 2020; Howard and Ruder, 2018; Liu et al., 2019; Lan et al., 2019; Raffel et al., 2020; Sun et al., 2019) retrains the model with new data. However, a key disadvantage is the large number of parameters that must be stored.

Efficient Fine-Tuning To alleviate the computational burden of fine-tuning, Adapter-BERT (Houlsby et al., 2019) proposes training only the adapter layers inserted into the model, while BitFit (Zaken et al., 2021) focuses on fine-tuning just the bias terms. Delta Tuning (Ding et al., 2022) explores parameter-efficient methods that adjust only a small portion of a model’s parameters. Low-Rank Adaptation methods (LoRA) (Hu et al., 2021)




Figure 3: **Comparison of Few-Shot Methods.** We compare CPT with baseline methods using the GPT-J model and the DBpedia dataset in few-shot settings, demonstrating its superior performance, particularly when handling a limited number of examples. Furthermore, our results highlight that context-based methods encounter memory constraints (indicated by dots) as the number of training examples increases beyond a certain threshold.

introduces a novel low-rank adaptation technique, where additional low-rank matrices are added to the weights during training. This allows the model to train only these matrices, reducing the number of trainable parameters significantly. VERA (Kopiczko et al., 2023) builds on LoRA by incorporating adaptive learning rates. Compacter (Karimi Mahabadi et al., 2021) leverages hyper-complex layers, and LoRA-Pro (Wang and Liang, 2024) further refines optimization. Despite these advancements, large models like GPT-3, which contain 175B parameters, require updating millions of parameters, such as 17.5M for LoRA.

Prompt Tuning (PT) Unlike fine-tuning methods, PT reduces the number of trainable parameters by introducing learnable tokens optimized while keeping the model’s weights frozen. (Lester et al., 2021) propose appending continuous prompts to the input and optimizing them, while P-tuning (Liu et al., 2023) and Prefix Tuning (Li and Liang, 2021) extend this concept by incorporating learnable tokens at intermediate layers. More recently, (Wang et al., 2023) introduced the idea of training a single prompt to be shared across multiple tasks. Although these methods significantly reduce the number of trainable parameters, they face challenges in few-shot learning (Gu et al., 2021) and provide limited interpretability for the learned continuous tokens (Ghosal et al., 2024; Khashabi et al., 2021; Deng et al., 2022).

In-Context Learning (ICL) In contrast to earlier methods, ICL (Brown et al., 2020) avoids optimization entirely. Instead, it concatenates task-specific examples before the input, allowing the model to learn a new task purely through observation, leveraging its pre-trained knowledge. Despite its advantages, ICL has limitations, often underperforming compared to optimization-based methods (Liu et al., 2022; Peng et al., 2023; Sun et al., 2023).

Instruction Prompt Tuning (IPT) IPT (Singhal et al., 2022) combines key elements of PT and ICL, utilizing learnable tokens that are optimized during training alongside static context tokens, similar to ICL. The concept of using both soft and hard prompts was previously introduced by PPT (Gu et al., 2021) and PTR (Han et al., 2022). Yet, IPT has struggled to consistently surpass PT in performance (Sun et al., 2023). While our method shares similarities with IPT, we focus on optimizing context tokens without introducing additional learnable tokens, and we are also leveraging context labels in the process. Another key difference lies in the

optimization process, where our loss includes a regularization term, and we employ projected gradient descent to ensure the output stays close to the user-supplied reliable input.

3 Our Method

3.1 Overfitting in Few-Shot Learning

In few-shot learning scenarios with limited data, the risk of overfitting in optimization-based methods is closely tied to the number of trainable parameters they introduce. For example, when working with LLaMA 3 8B, methods such as full fine-tuning, LoRA, and PT involve updating approximately 8B, 4.2M, and 32K parameters, respectively. This demonstrates that even the most parameter-efficient optimization-based methods still require training a significant number of parameters, which can pose challenges when data is limited.

Our analysis, presented in fig. 1, confirms that all optimization-based methods, including LoRA and PT, exhibit a train-test loss gap, highlighting overfitting in low-data regimes where the number of examples ranges from 2 to 20. Despite variations in the number of trainable parameters, these methods struggle to generalize effectively when data is limited. These findings emphasize the need for approaches like CPT, which is carefully designed to balance flexibility and generalization, effectively mitigating overfitting challenges in few-shot learning scenarios.

3.2 Input Preparation

Our method takes as input a few-shot classification dataset containing N examples. Each example consists of a pairing of x (an instruction) and y (a label). We embed (x, y) using input, output, and separation templates, converting them into readable text that LLMs better understand, as done in ICL (Brown et al., 2020). The input and output templates, denoted T_i and T_o , along with separators S_{intra} and S_{inter} , are provided in appendix E. To embed a single example (x, y) using the template, we concatenate the input x embedded in T_i with S_{intra} , followed by the output y embedded in T_o , and finally S_{inter} , resulting in $X_{\text{Emb}_i} = [T_i(x_i), S_{\text{intra}}, T_o(y_i), S_{\text{inter}}]$. To generate the complete context, we concatenate all X_{Emb_i} , forming $X_{\text{Context}} = [X_{\text{Emb}_i}]_{i=1}^N$. To construct a complete training example, we randomly select an embedded example from the training set X_{Emb_i} , and concatenate it after the context, result-

ing $X_{\text{Train}_i} = [X_{\text{Context}}, X_{\text{Emb}_i}]$, which is then fed into the LLM. This process is also visualized in fig. 4, with additional concrete examples provided in appendix G.

Above, we described how we construct a training example X_{Train_i} , as a text sequence. However, before feeding it into the model, we must process the text through a tokenizer, which splits the text into tokens and returns an embedding vector for each token. Each example contains six types of tokens: input, input template, intra-separator, output, output template, and inter-separator. For simplicity, we ignore the separators and the fact that each part usually contains multiple tokens. For each training example i and its sub-example k , we focus on four token types: $t_{I_i}^{(k)}$, $t_{IT_i}^{(k)}$, $t_{O_i}^{(k)}$, $t_{OT_i}^{(k)}$, which represent the input, input template, output, and output template, respectively. Each training example i consists of $N + 1$ sub-examples, N sub-examples in the context and one training sub-example at the end.

3.3 Optimization

In this section, we discuss the optimization process of our method, which draws inspiration from Adversarial Attacks (AT) (Madry et al., 2017). The AT process typically consists of two key components: optimization and restriction. First, an attacker modifies an image to induce misclassification; second, the attack constrains its changes to evade detection. Inspired by this approach, our method follows a similar structure, consisting of two key phases: optimization, which encompasses the loss design outlined in section 3.3.1, and regulating token updates, as detailed in section 3.3.2.

3.3.1 Loss Design

The optimization process aims to refine the input embeddings to enhance classification performance. To achieve this, we introduce a novel loss function for each training example X_{Train_i} , which incorporates all the context sub-example labels X_{Context} . More formally, the loss compares the model’s predicted values $\hat{t}_{O_i}^{(k)}$ to the ground truth tokens $t_{O_i}^{(k)}$ for all $k \in [1, N]$. These target tokens and their corresponding predictions, as illustrated in eq. (1), form the basis of our optimization objective.

$$L_{\text{Context}_i} = \sum_{k=1}^N \omega_k \cdot \text{CrossEntropy}(\hat{t}_{O_i}^{(k)}, t_{O_i}^{(k)}) \quad (1)$$

In addition to L_{Context_i} , we also apply the standard loss on the training sub-example in eq. (2).

$$L_{\text{Train}_i} = \text{CrossEntropy}(\hat{t}_{O_i}^{(N+1)}, t_{O_i}^{(N+1)}) \quad (2)$$

Lastly, we sum both losses to create the final loss $L_i = L_{\text{Context}_i} + L_{\text{Train}_i}$, where L_{Context_i} can be thought of as a regularization for the standard loss L_{Train_i} .

As explained in section 3.2, each training example X_{Train_i} contains $N + 1$ sub-labels, from N sub-examples in the context and one training sub-example. However, not all sub-examples should be weighted equally. For instance, the last sub-example is more important as it is located in the location of the test examples. Additionally, sub-examples closer to the end of the context carry more importance (Zhao et al., 2021). Thus, we apply exponential loss weight decay starting from the end of the context and decaying towards the beginning, while keeping L_{Train_i} unchanged. Formally, each sub-example k is multiplied by γ^j , where $j = N + 1 - k$. For example, the last sub-example is multiplied by γ^1 , and the second-to-last by γ^2 , and so on. The decay is shown in eq. (1) as ω_k .

3.3.2 Controlled Token Embedding Optimization

As mentioned in section 3.3.1, we utilize all the labels within each training example X_{Train_i} to optimize the context tokens. However, within the context X_{Context_i} , some tokens serve as labels and are crucial for the optimization process. Therefore, we keep these label tokens fixed, as they carry valuable information that acts as a regularization component in the loss term, as explained in section 3.3.1. The remaining context tokens are updated through the optimization process in a carefully controlled manner, as detailed in the following section.

The controlled optimization process is an effective strategy for improving generalization and is commonly addressed through techniques such as the Adam optimizer (Kingma, 2014), which limits the magnitude of model weight updates. In our method, all updates are applied to the context; thus, constraining these changes prevents the model from overfitting to the provided training examples and instead promotes better generalization. Additionally, maintaining the context’s proximity to the user-provided examples enhances interpretability—similar to ICL—allowing the model to lever-




Figure 4: **Constructing a Training Example with CPT**. This figure illustrates how CPT builds the second training example, X_{Train_2} , in a dataset containing two examples. We begin by concatenating the embedded input-output templates, $[X_{\text{Emb}_i}]_{i=1}^2$, to create the context, X_{Context} . The embedding of the second example, X_{Emb_2} , is then appended to X_{Context} , forming the complete input X_{Train_2} . This input is passed through the frozen LLM, and the loss is computed using all labels in X_{Train_2} , which include both the context and training labels. Importantly, while the context is updated during this process, its labels remain unchanged.

age meaningful examples without introducing excessive modifications.

To achieve this, we employ projected gradient descent (PGD), a widely used technique in adversarial attacks, to constrain token embedding updates within an ℓ_2 norm of size ϵ after each optimization step. This ensures controlled adjustments while preserving the integrity of the original context. Further details are provided in appendix H. As the modification norm ϵ decreases, our method gradually converges to ICL, which inherently exhibits robustness against overfitting. This controlled optimization strikes a balance between flexibility and stability, enabling the model to refine the context while mitigating the risk of overfitting.

4 Experimental Setup

In this section, we provide details regarding the datasets, models, baselines, and evaluation used in our experiments. Implementation details are provided in appendix F.

4.1 Datasets

In this work, we focus on a classification task and select a variety of datasets to ensure robust conclusions across different task types. We include SST-2 (Socher et al., 2013) for sentiment analysis, AG News (Zhang et al., 2015b) for news classification, DBpedia (Zhang et al., 2015a) for ontology classification,

and TREC (Li and Roth, 2002) for question classification. These datasets represent a diverse range of natural language classification tasks, include different number of classification classes, allowing us to evaluate our method comprehensively. More details are provided in appendix D.

4.2 Models

We use models of varying sizes and quality to ensure robust evaluation and conclusions. For the relatively small model, we use BLOOM1.7B (Scao et al., 2022), while for larger models, we opt for GPT-J6B (Wang and Komatsuzaki, 2021) and Llama3 8B (AI@Meta, 2024). The GPT-J model is noted for its robust performance, while Llama3 is currently among the leading models in the field.

4.3 Baselines

We compare our method to several groups of few-shot learning techniques. In the first group, we include LoRA (Hu et al., 2021), one of the leading efficient fine-tuning methods. Additionally, we compare against several prompt-tuning approaches, including Prompt Tuning (PT) (Lester et al., 2021), Prefix Tuning (Li and Liang, 2021), and Instruction Prompt Tuning (IPT) (Singhal et al., 2022). Finally, we compare our method to In-Context Learning (ICL) (Brown et al., 2020).

For some of the few-shot methods, we introduce an alternative version that incorporates instructions,




Figure 5: **Accuracy and Standard Deviation** Comparison of accuracy and standard deviation between CPT and baselines, evaluated with 4-shot on GPT-J model. The black bars represent the mean std across different templates, while the blue bars represent the mean std across different seeds.

as indicated in table 1 with a †. Instead of initializing the learnable tokens randomly, we initialize them with instructions specified in appendix C. We apply instructions to PT, IPT, and our method, reporting results for both random and instruction-based prompt initialization. An example illustrating how inputs are constructed with and without † is provided in appendix G.

4.4 Evaluation

We evaluate each model and dataset using three different numbers of training samples: 2, 4, and 6. For each configuration, the reported results are averaged accuracy over 30 experiments, consisting of 10 randomly sampled templates and 3 different random seeds, with the templates described in appendix E. By utilizing randomized seeds, we ensure variation in the selection of training examples. This extensive setup is crucial for achieving a comprehensive and robust evaluation, especially given that these methods are known to be highly sensitive to the selection of training examples and templates (Voronov et al., 2024; Zhao et al., 2021). Further evaluation details can be found in appendix B.

5 Results

5.1 Main Results

In table 1, we demonstrate that CPT convincingly performs better than the baselines in most cases, with particularly pronounced gains in harder tasks. Furthermore, CPT’s performance becomes more efficient and effective as the models grow stronger, such as with Llama3.

Performance on Challenging Tasks CPT demonstrates improvements across various datasets, with more pronounced gains in tasks we define as harder based on two factors: the number of shots and the number of classes. As illustrated

in table 1, task difficulty increases with fewer shots and more classes. For example, on the DBpedia dataset, which has 14 classes, decreasing the shots from 6 to 4 widens the performance gap between CPT and the baselines from (3, 6, 1) to (11, 10, 3) across the models: BLOOM, GPT-J, and Llama3.

Decisive Advantage with Powerful Models

The strength of the model plays a significant role in performance. As the model becomes better, CPT’s advantage becomes more pronounced across all datasets and shot settings. For instance, Llama3 consistently outperforms other baselines across all datasets, except in one case where results are comparable. With GPT-J, a slightly older model, the results are lower in two instances, with one comparable outcome, both on SST-2, the easier task as previously discussed. When comparing with BLOOM, the weakest model in our comparison, we observe lower performance on two occasions, specifically on the two easier datasets.

5.2 Standard Deviation

Standard deviation (std) plays a crucial role in few-shot learning due to the sensitivity of these methods to both the training examples and the chosen template (Zhao et al., 2021; Voronov et al., 2024). In fig. 5, we present accuracy along with two types of std bars: black bars represent the mean std across different templates, while blue bars represent the mean std across different seeds. We demonstrate that CPT significantly improves accuracy across various models and datasets in a statistically significant manner. More information is presented in appendix A.

Our method’s standard deviation performs equivalently to other methods in most cases, while in certain cases, such as with DBpedia, CPT exhibits both higher accuracy and lower std, reinforcing its

Dataset	Method	BLOOM 1.7B			GPT-J 6B			Llama3 8B		
		2	4	6	2	4	6	2	4	6
SST-2	Prefix	47.80	47.33	49.00	52.23	52.50	52.87	—	—	—
	ICL	50.53	60.83	61.87	50.57	67.47	77.47	76.43	80.63	83.10
	PT†	64.97	65.07	65.07	57.10	52.93	55.70	72.97	73.47	84.57
	PT	56.03	56.90	58.33	64.07	64.37	64.60	64.27	65.70	67.03
	IPT†	58.50	61.83	62.80	51.50	83.20	84.80	86.90	88.03	94.40
	IPT	48.50	58.80	61.87	48.13	82.27	87.17	57.20	87.40	90.43
	LoRA	66.40	66.93	66.90	69.80	71.53	73.17	68.73	71.27	83.97
	CPT†	59.53	72.40	74.83	52.53	82.03	88.07	92.73	95.07	96.40
	CPT	50.77	70.70	74.10	50.53	82.90	88.03	83.83	96.30	96.50
AG News	Prefix	24.87	25.35	26.02	32.32	33.33	46.08	—	—	—
	ICL	35.12	34.28	42.48	66.73	62.38	69.57	79.38	82.32	85.27
	PT†	28.67	30.73	41.17	37.85	44.85	62.92	59.60	57.02	68.02
	PT	33.57	36.98	56.08	56.85	56.13	75.10	69.32	67.92	69.33
	IPT†	36.95	31.90	42.93	67.02	63.00	74.85	82.93	84.45	85.08
	IPT	38.77	38.20	47.78	66.02	63.92	74.00	80.52	76.30	80.98
	LoRA	29.50	30.80	33.98	56.12	56.03	72.55	70.62	74.97	73.70
	CPT†	33.68	33.13	41.10	71.35	68.73	75.68	83.17	84.28	84.67
	CPT	40.85	44.48	50.40	74.80	68.62	76.22	83.78	81.92	85.43
DBpedia	Prefix	19.76	19.74	23.65	13.25	16.43	24.94	—	—	—
	ICL	48.20	51.40	55.17	50.87	62.46	70.76	71.66	72.44	79.93
	PT†	24.90	26.32	34.75	21.01	22.12	37.44	55.30	57.21	66.26
	PT	46.71	41.94	45.93	23.39	29.69	40.53	55.81	52.72	55.02
	IPT†	33.28	40.36	45.85	47.10	67.60	75.09	81.10	87.69	92.06
	IPT	48.09	54.60	70.57	52.86	67.27	70.73	72.92	76.11	78.44
	LoRA	43.30	41.13	41.18	30.15	28.02	41.50	54.24	59.50	63.21
	CPT†	33.80	48.13	51.18	53.20	77.30	81.00	84.23	90.33	93.08
	CPT	58.85	65.78	73.55	68.29	75.07	77.65	77.38	78.49	82.42
TREC	Prefix	19.10	24.49	29.92	30.76	30.04	27.87	—	—	—
	ICL	33.54	33.33	28.53	28.94	35.14	32.49	35.32	42.48	40.34
	PT†	30.91	33.70	39.31	29.02	34.66	43.89	43.42	48.81	51.73
	PT	32.18	32.26	35.69	31.16	32.79	37.86	32.77	33.98	33.83
	IPT†	27.83	36.64	42.92	31.04	43.12	43.09	51.72	62.14	65.13
	IPT	32.37	36.59	42.60	29.59	38.90	40.38	36.94	45.62	52.08
	LoRA	34.07	33.22	33.50	34.17	33.73	37.63	31.21	33.21	36.36
	CPT†	29.72	35.64	45.38	33.39	44.20	45.83	57.26	67.00	69.29
	CPT	35.68	41.79	45.16	35.37	44.66	42.71	45.12	57.54	60.18

Table 1: **Baseline Comparisons** Mean accuracy of various methods and our CPT, across several models and datasets. Evaluations are conducted using 2, 4, and 6 shots.

robustness in complex tasks. However, the sensitivity of our method does not follow a clear pattern across random seeds or templates. For instance, while randomness in templates and training examples has an equal influence on std in DBpedia and TREC, SST-2 shows a higher std for template randomness, and AG News is more sensitive to variations in training examples.

5.3 Ablations

Our ablation studies aim to dissect the contributions of individual components in CPT, highlighting the elements that drive its performance improvements across few-shot learning tasks, as shown in table 2. As shown, the loss design and the projections are the most important component of our method. Further ablation experiments can be found in appendix I.

Loss Design Different options for the loss design

are specified under “Loss Tokens”, with three configurations: using only the training label, using the training label plus one random context label, and using the training label plus all context labels. The latter outperforms the training-only configuration by 11%, 12%, 10% for 2, 4 and 6 shots.

Effect of Projection Magnitude The ablation study on projection magnitude is specified under “Input ϵ ” and “Format ϵ ”, which define the allowable deviation from the original values for input tokens and format tokens, respectively. The results demonstrate that both excessively small changes (leading to convergence toward ICL) and overly large norms (failing to limit overfitting) are suboptimal, emphasizing the importance of selecting an appropriate projection magnitude.

Loss Weighting We evaluated the impact of different loss weighting strategies and propose three options: (1) *Mean*, which applies uniform weight-

Loss Tokens	Loss Weighting	Projection Type	Input ϵ	Format ϵ	Updated Tokens	Mask Training Example	Number of Training Examples		
							2	4	6
Train Example							58.09	61.54	66.69
Train Example & 1 Random	Decay 0.95	Token-Wise	0.1	0.1	Input & Format	✗	69.48	72.08	76.80
Train Example & All Context							69.54	73.03	76.58
Train Example & All Context	Mean	Token-Wise	0.1	0.1	Input & Format	✗	69.62	72.91	76.49
	Equal 1						69.07	72.82	76.23
	Equal 10						69.35	71.01	75.11
	Decay 0.99						69.59	72.97	76.43
	Decay 0.95						69.54	73.03	76.58
	Decay 0.5						69.60	72.39	76.44
Train Example & All Context	Decay 0.95	All-Tokens	0.001	-	Input & Format	✗	51.52	63.41	71.50
			0.01	-			56.37	68.12	73.66
			0.1	-			69.51	72.64	76.06
			1.0	-			63.11	64.78	71.94
Train Example & All Context	Decay 0.95	Token-Wise	0.01	0.1	Input & Format	✗	65.61	70.12	75.63
			0.1	0.1			69.54	73.03	76.58
			1.0	0.1			65.29	66.30	73.63
			0.1	0.01			69.53	73.55	76.55
			0.1	1.0			68.27	71.91	68.27
Train Example & All Context	Decay 0.95	Token-Wise	0.1	0.1	Input Masks	✗	69.47	74.13	76.63
					Input & Format		63.74	69.21	74.91
					Input & Format		69.54	73.03	76.58
Train Example & All Context	Decay 0.95	Token-Wise	0.1	0.1	Input & Format	✓	67.55	64.26	68.58

Table 2: **Ablation Study** We present the mean accuracy for various ablations using the GPT-J model and the DBpedia dataset, including loss tokens (train example, random, or all context), loss weighting (decay and mean), projection type (token-wise or all-tokens), epsilon values for input and format, updated tokens (input, format, masks), and masking of the training example.

ing across all labels; (2) *Equal*, which assigns equal weight to the training label loss and the context label losses, with an optional scaling factor applied to the training loss (*e.g.*, 1, 10); and (3) *Decay*, which exponentially reduces the weight of context labels further from the training example, with the decay factor specified (*e.g.*, 0.99, 0.95, 0.5).

Projection Type: Token-wise vs. All-Tokens

We evaluated the “All-Tokens” projection approach, which applies the projection to the entire context collectively rather than processing it token-by-token. Our results indicate that the token-wise approach is preferable, as it provides stronger regularization by limiting each token individually rather than the context as a whole, resulting in better performance.

Updated Tokens Under “Updated Tokens”, we explored modifying only specific parts of the context to determine if certain components are more critical for updates than others. Our results indicate that allowing changes to both the input and format tokens yields better performance, provided these changes are constrained using the projection limitation.

Mask Training We also experimented with “Mask Training,” where the training example was masked from the context to prevent the model from simply copying the answer. In our setup, the training example appears both in the context (along with the correct answer) and as an additional concatenated example at the end. Masking the training example from the context and removing this duplication seemed like a plausible strategy to improve

generalization. However, this approach did not lead to any performance improvements.

6 Discussions

In this work, we identify overfitting as the primary reason for the underperformance of optimization-based methods in few-shot learning scenarios, substantiated by empirical evidence. To address this challenge, we propose CPT, an optimization-based method that effectively mitigates overfitting. Our results demonstrate that CPT consistently outperforms existing baselines across diverse datasets, models, and experimental setups.

Beyond its direct contributions to few-shot learning, CPT highlights the critical importance of balancing optimization flexibility and regularization in data-scarce scenarios. The insights from this work can inspire the development of parameter-efficient, robust, and interpretable approaches for a range of machine learning challenges, including transfer learning, domain adaptation, and fine-tuning in resource-constrained environments.

Limitation & Future Work The computational cost associated with the iterative optimization of context embeddings is significant compared to ICL. Additionally, similar to ICL and IPT, CPT is limited in the number of examples it can handle, as memory consumption scales with context length. In contrast, traditional methods are better suited for larger datasets. Future work could explore more efficient optimization strategies to reduce computational overhead and improve scalability.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. 2018. Synthesizing robust adversarial examples. In *International conference on machine learning*, pages 284–293. PMLR.
- Tsachi Blau, Roy Ganz, Chaim Baskin, Michael Elad, and Alex Bronstein. 2023. Classifier robustness enhancement via test-time transformation. *arXiv preprint arXiv:2303.15409*.
- Tsachi Blau, Roy Ganz, Bahjat Kavar, Alex Bronstein, and Michael Elad. 2022. Threat model-agnostic adversarial defense using diffusion models. *arXiv preprint arXiv:2207.08089*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Nicholas Carlini and David Wagner. 2017. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 3–14.
- Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P Xing, and Zhiting Hu. 2022. Rlprompt: Optimizing discrete text prompts with reinforcement learning. *arXiv preprint arXiv:2205.12548*.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. 2022. Delta tuning: A comprehensive study of parameter efficient methods for pre-trained language models. *arXiv preprint arXiv:2203.06904*.
- Soumya Suvra Ghosal, Samyadeep Basu, Soheil Feizi, and Dinesh Manocha. 2024. Intcoop: Interpretability-aware vision-language prompt tuning. *arXiv preprint arXiv:2406.13683*.
- Sven Gowal, Chongli Qin, Jonathan Uesato, Timothy Mann, and Pushmeet Kohli. 2020. Uncovering the limits of adversarial training against norm-bounded adversarial examples. *arXiv preprint arXiv:2010.03593*.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. 2022. Ptr: Prompt tuning with rules for text classification. *AI Open*, 3:182–192.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- 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](#). *Preprint*, arXiv:2106.09685.
- Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. 2021. Compacter: Efficient low-rank hypercomplex adapter layers. *Advances in Neural Information Processing Systems*, 34:1022–1035.
- Daniel Khashabi, Shane Lyu, Sewon Min, Lianhui Qin, Kyle Richardson, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, Sameer Singh, et al. 2021. Prompt waywardness: The curious case of discretized interpretation of continuous prompts. *arXiv preprint arXiv:2112.08348*.
- Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Dawid J. Kopiczko, Tijmen Blankevoort, and Yuki M. Asano. 2023. [Vera: Vector-based random matrix adaptation](#). *Preprint*, arXiv:2310.11454.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). *Preprint*, arXiv:2104.08691.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Xin Li and Dan Roth. 2002. [Learning question classifiers](#). In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin A Raffel. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *Advances in Neural Information Processing Systems*, 35:1950–1965.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. Gpt understands, too. *AI Open*.
- 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. *arXiv preprint arXiv:1907.11692*.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2017. Towards deep learning models resistant to adversarial attacks. <i>arXiv preprint arXiv:1706.06083</i> .	721
Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. 2022. Peft: State-of-the-art parameter-efficient fine-tuning methods. https://github.com/huggingface/peft .	722
Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? <i>arXiv preprint arXiv:2202.12837</i> .	723
Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. 2023. Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation. <i>arXiv preprint arXiv:2305.16938</i> .	724
Hao Peng, Xiaozhi Wang, Jianhui Chen, Weikai Li, Yunjia Qi, Zimu Wang, Zhili Wu, Kaisheng Zeng, Bin Xu, Lei Hou, et al. 2023. When does in-context learning fall short and why? a study on specification-heavy tasks. <i>arXiv preprint arXiv:2311.08993</i> .	725
Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. <i>OpenAI blog</i> , 1(8):9.	726
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>Journal of machine learning research</i> , 21(140):1–67.	727
Nir Ratner, Yoav Levine, Yonatan Belinkov, Ori Ram, Inbal Magar, Omri Abend, Ehud Karpas, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2022. Parallel context windows for large language models. <i>arXiv preprint arXiv:2212.10947</i> .	728
Teven Le Scao, Angela Fan, Christopher Akiki, and et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. <i>arXiv preprint arXiv:2211.05100</i> .	729
Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2022. Large language models encode clinical knowledge. <i>arXiv preprint arXiv:2212.13138</i> .	730
Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. <i>Recursive deep models for semantic compositionality over a sentiment treebank</i> . In <i>Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing</i> , pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.	731
Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune bert for text classification? In <i>Chinese computational linguistics: 18th China national conference, CCL 2019, Kunming, China, October 18–20, 2019, proceedings 18</i> , pages 194–206. Springer.	732
Simeng Sun, Yang Liu, Dan Iter, Chenguang Zhu, and Mohit Iyyer. 2023. How does in-context learning help prompt tuning? <i>arXiv preprint arXiv:2302.11521</i> .	733
Anton Voronov, Lena Wolf, and Max Ryabinin. 2024. Mind your format: Towards consistent evaluation of in-context learning improvements. <i>arXiv preprint arXiv:2401.06766</i> .	734
Ben Wang and Aran Komatsuzaki. 2021. Gpt-j-6b: A 6 billion parameter autoregressive language model. https://github.com/kingoflolz/mesh-transformer-jax . Accessed: 2024-05-26.	735
Zhen Wang, Rameswar Panda, Leonid Karlinsky, Rogério Feris, Huan Sun, and Yoon Kim. 2023. Multitask prompt tuning enables parameter-efficient transfer learning. <i>arXiv preprint arXiv:2303.02861</i> .	736
Zhengbo Wang and Jian Liang. 2024. Lora-pro: Are low-rank adapters properly optimized? <i>arXiv preprint arXiv:2407.18242</i> .	737
Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. <i>Transformers: State-of-the-art natural language processing</i> .	738
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> .	739
Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015a. <i>Character-level convolutional networks for text classification</i> . In <i>Advances in Neural Information Processing Systems</i> , volume 28. Curran Associates, Inc.	740
Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015b. <i>Character-level convolutional networks for text classification</i> . In <i>NIPS</i> .	741
Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In <i>International conference on machine learning</i> , pages 12697–12706. PMLR.	742
A Standard Deviation	743
In table 3 we present the standard deviations (STD) corresponding to the main results shown in Table 1. For each experiment, we display three STD values,	744

Dataset	Method	BLOOM 1.7B						Model GPT-J 6B						Llama3 8B					
		2	4	6	2	4	6	2	4	6	2	4	6	2	4	6	2	4	6
SST-2	Prefix Tuning	00.5/00.4/00.1	03.1/02.7/01.9	03.1/02.6/02.1	00.8/00.6/00.2	02.3/01.5/00.7	05.8/04.3/03.9	—	—	—	—	—	—	—	—	—	—	—	—
	ICL	04.3/04.0/01.5	12.6/08.6/09.4	14.9/10.6/09.7	05.5/04.0/03.0	14.1/12.6/08.2	13.1/09.9/09.9	13.2/12.7/06.1	15.7/11.9/10.2	13.7/12.2/06.5	—	—	—	—	—	—	—	—	—
	PT†	07.6/07.6/00.4	08.2/08.2/00.7	08.1/08.1/00.6	06.8/06.5/02.3	07.8/06.9/04.3	09.5/09.0/04.9	16.6/16.5/01.0	17.1/17.1/01.6	12.7/10.9/05.7	—	—	—	—	—	—	—	—	—
	PT	08.6/08.5/01.7	08.9/08.6/02.4	08.7/08.4/02.5	07.7/07.6/01.3	07.0/07.0/00.9	07.4/07.4/01.0	06.4/06.1/03.1	06.8/06.6/02.8	07.0/06.5/03.7	—	—	—	—	—	—	—	—	—
	IPT†	12.3/12.3/03.4	14.1/11.4/09.8	15.3/10.1/12.6	05.8/05.1/02.5	11.1/08.0/08.3	12.3/12.3/02.0	08.7/06.8/05.6	12.5/11.2/07.3	02.7/02.6/01.6	—	—	—	—	—	—	—	—	—
	IPT	02.0/01.5/00.4	11.6/08.7/08.2	14.8/09.0/11.8	00.7/00.4/00.2	13.0/08.1/10.0	07.2/05.7/04.7	13.7/13.7/03.9	10.6/09.6/05.5	12.0/10.3/05.3	—	—	—	—	—	—	—	—	—
	LoRA	06.9/06.9/00.2	06.5/06.5/00.5	06.4/06.3/00.5	09.5/09.5/00.8	10.2/10.2/01.2	10.0/09.9/01.4	11.4/11.4/01.4	15.0/14.6/07.6	12.1/11.9/07.2	—	—	—	—	—	—	—	—	—
	CPT†	12.3/12.3/03.6	13.8/09.8/10.5	15.4/14.2/09.6	07.6/06.7/03.2	11.1/09.1/07.0	07.0/06.2/04.1	05.0/04.3/02.5	04.1/03.0/02.3	02.0/01.7/01.2	—	—	—	—	—	—	—	—	—
	CPT	07.7/05.3/02.9	12.0/11.1/06.9	12.9/10.5/09.7	05.5/03.9/02.9	12.7/12.0/05.9	10.7/08.9/05.1	13.2/10.9/08.5	01.6/01.5/01.0	01.3/01.2/01.0	—	—	—	—	—	—	—	—	—
	Prefix Tuning	01.7/01.7/00.6	05.8/02.9/05.2	05.3/03.9/03.9	05.9/05.9/00.5	06.2/06.1/01.4	12.7/09.4/08.2	—	—	—	—	—	—	—	—	—	—	—	—
AG News	ICL	10.5/06.8/08.7	11.7/10.1/06.4	12.2/11.1/06.0	10.0/08.9/05.2	13.3/10.6/09.2	10.4/09.4/05.0	08.8/03.2/08.1	03.2/03.0/02.3	03.1/02.7/02.4	—	—	—	—	—	—	—	—	—
	PT†	05.2/04.1/04.1	06.0/04.3/04.4	10.9/10.5/07.3	16.2/16.1/00.9	13.7/11.5/09.5	13.6/11.9/07.8	11.9/11.8/06.1	10.9/10.3/05.8	10.5/09.0/06.5	—	—	—	—	—	—	—	—	—
	PT	07.9/04.6/06.9	08.4/06.8/05.7	09.9/09.3/05.2	12.3/11.5/05.3	10.8/06.7/08.6	07.0/02.5/06.8	15.0/15.0/01.5	15.4/15.2/02.7	12.4/12.3/02.2	—	—	—	—	—	—	—	—	—
	IPT†	11.7/09.2/08.4	07.5/05.5/05.0	15.0/09.4/12.4	11.0/08.3/07.9	07.1/03.1/06.5	07.6/03.5/06.8	02.7/02.5/01.6	03.5/02.8/02.2	03.4/02.9/02.5	—	—	—	—	—	—	—	—	—
	IPT	12.0/08.6/08.9	10.6/09.1/07.2	12.0/09.4/07.9	11.1/10.0/06.3	08.8/04.6/07.6	07.4/04.6/05.7	03.6/03.1/01.8	08.1/03.2/07.6	05.2/02.2/04.7	—	—	—	—	—	—	—	—	—
	LoRA	03.2/03.2/00.5	06.8/03.3/06.1	04.6/04.4/02.7	09.0/09.0/01.3	09.4/09.3/01.8	08.8/05.3/07.3	15.0/15.0/01.5	15.4/15.2/02.7	12.4/12.3/02.2	—	—	—	—	—	—	—	—	—
	CPT†	07.2/05.7/05.4	06.0/03.4/04.8	11.9/08.5/09.2	09.6/07.2/07.1	09.0/03.6/08.5	09.2/04.4/08.1	03.1/02.5/02.2	02.9/02.6/02.4	03.4/01.8/03.2	—	—	—	—	—	—	—	—	—
	CPT	12.8/08.7/10.3	12.2/07.8/10.3	11.3/08.9/07.1	08.6/05.4/07.4	09.1/03.8/08.5	07.2/03.7/06.2	03.3/02.6/02.3	04.3/03.6/02.8	02.9/02.3/02.3	—	—	—	—	—	—	—	—	—
	Prefix Tuning	03.5/03.4/01.9	06.4/03.1/05.7	08.7/03.6/08.1	02.3/02.3/01.7	04.5/02.7/03.6	07.9/04.9/06.2	—	—	—	—	—	—	—	—	—	—	—	—
	ICL	24.1/23.3/06.9	25.8/23.5/08.9	23.9/23.6/06.0	16.6/16.3/05.9	15.7/13.1/08.1	06.7/05.8/04.0	07.7/06.4/06.2	06.8/02.6/06.5	04.2/02.3/04.0	—	—	—	—	—	—	—	—	—
DBpedia	PT†	15.6/08.5/13.2	09.7/09.6/01.7	07.1/04.6/05.8	10.3/10.3/00.9	06.3/05.8/04.2	06.3/05.8/04.2	19.5/17.3/11.0	15.7/12.7/09.0	15.4/13.7/08.2	—	—	—	—	—	—	—	—	—
	PT	11.2/11.1/04.2	10.8/10.8/01.3	13.0/12.4/05.9	09.9/08.2/06.4	11.3/07.1/09.1	08.9/05.0/07.5	11.9/11.7/04.2	15.7/15.3/02.9	13.5/13.4/01.5	—	—	—	—	—	—	—	—	—
	IPT†	25.6/21.2/14.7	24.3/22.6/08.9	27.3/26.2/08.5	16.4/15.5/05.7	11.0/09.7/06.2	06.8/05.3/05.2	05.3/04.3/03.0	04.5/03.8/02.7	04.5/04.1/01.9	—	—	—	—	—	—	—	—	—
	IPT	26.2/25.1/07.5	25.0/20.3/11.9	07.6/07.0/02.9	12.2/11.2/06.0	09.6/06.0/07.2	05.4/03.7/04.1	09.7/08.4/04.6	06.0/03.6/05.1	05.6/03.1/05.3	—	—	—	—	—	—	—	—	—
	LoRA	11.4/11.0/03.1	11.6/11.6/00.3	11.7/11.7/00.4	11.6/10.2/04.9	13.0/06.0/11.6	09.8/06.0/07.9	13.1/13.0/01.7	14.3/14.2/02.2	13.7/13.7/01.5	—	—	—	—	—	—	—	—	—
	CPT†	23.2/14.5/18.0	12.0/10.1/07.1	22.1/20.1/10.5	15.6/08.6/14.3	06.5/05.0/04.5	06.0/03.2/05.2	06.2/05.6/02.7	03.8/03.4/02.4	02.3/02.2/01.7	—	—	—	—	—	—	—	—	—
	CPT	15.5/13.4/06.5	11.8/08.8/07.1	04.7/04.0/02.5	10.9/08.2/05.5	05.0/03.8/03.5	03.9/03.0/03.0	06.1/05.0/04.8	04.3/03.5/03.0	04.3/02.6/03.5	—	—	—	—	—	—	—	—	—
	Prefix Tuning	06.7/00.8/06.6	06.4/02.9/06.0	07.0/04.6/06.0	03.1/02.0/02.5	05.9/02.6/05.1	03.7/03.6/00.8	—	—	—	—	—	—	—	—	—	—	—	—
	ICL	11.0/07.2/08.2	10.5/08.2/06.8	13.8/09.0/09.1	08.9/05.9/06.8	11.0/08.2/07.8	12.6/08.3/09.4	08.6/05.6/06.3	14.2/07.9/12.0	13.2/08.3/10.7	—	—	—	—	—	—	—	—	—
	PT†	05.9/04.4/03.7	07.5/06.7/04.5	11.2/08.7/07.8	05.5/04.3/06.8	06.2/06.0/04.3	13.5/08.2/11.7	09.5/05.7/08.3	11.3/06.1/09.4	10.3/08.4/08.1	—	—	—	—	—	—	—	—	—
TREC	PT	03.8/03.4/01.5	08.2/07.3/06.7	11.2/08.6/09.1	04.0/04.0/00.9	08.1/07.5/05.7	09.7/08.2/07.8	05.0/05.0/01.5	04.5/04.5/02.5	03.8/03.8/02.0	—	—	—	—	—	—	—	—	—
	IPT†	05.5/03.6/04.0	10.1/09.1/07.1	16.8/08.1/15.5	06.8/05.3/04.2	07.7/05.1/05.9	14.0/07.8/11.9	13.6/06.3/12.3	09.7/05.9/08.6	07.2/05.4/05.5	—	—	—	—	—	—	—	—	—
	IPT	10.5/07.1/07.8	06.1/05.9/05.0	14.1/07.4/12.8	09.7/05.4/08.0	09.3/05.7/07.5	13.0/03.8/12.5	12.1/08.5/07.8	11.5/09.1/08.0	14.0/05.4/13.3	—	—	—	—	—	—	—	—	—
	LoRA	03.9/03.9/01.0	04.0/04.0/00.3	04.1/04.1/00.4	02.5/02.5/00.4	03.6/03.5/02.2	11.9/07.2/10.4	03.3/03.3/01.0	03.5/02.7/02.5	16.5/08.1/15.4	—	—	—	—	—	—	—	—	—
	CPT†	08.0/05.8/05.4	07.7/06.9/06.3	09.9/07.0/07.9	08.5/05.7/06.4	12.9/08.3/10.6	11.2/08.2/09.0	13.1/06.9/11.6	09.8/03.7/09.2	05.0/04.1/03.4	—	—	—	—	—	—	—	—	—
	CPT	09.1/05.2/07.3	07.9/07.2/05.6	12.9/07.0/10.8	07.4/04.1/06.1	08.7/06.2/06.9	08.6/05.5/07.3	16.8/08.4/14.5	07.4/06.5/05.8	07.9/05.6/06.5	—	—	—	—	—	—	—	—	—

Table 3: **Standard Deviation Analysis** Standard deviations (STD) corresponding to Table 1. Each experiment shows three STD values separated by a backslash: (1) STD over 30 experiments with 10 random templates and 3 seeds, (2) mean STD over templates, and (3) mean STD over seeds.

separated by a backslash. These values represent the variability in the results across different configurations:

1. The first value shows the standard deviation over 30 experiments, which includes 10 random templates and 3 seeds that determine the training examples.
2. The second value provides the mean of the standard deviation over the templates, the standard deviation across 10 templates, and the mean of the standard deviation across 3 seeds.
3. The third value presents the mean standard deviation over the seeds, the standard deviation over 3 seeds, and the mean over 10 templates.

This detailed breakdown of standard deviations allows for a more thorough understanding of the variability in model performance across different templates and seeds.

B Evaluation Details

All the graphs and ablation studies were conducted and evaluated using the DBpedia dataset with the GPT-J model. This setup was chosen due to the diversity of the DBpedia dataset, which includes a broad range of categories and entities, making it an ideal candidate for comprehensive evaluation. The use of GPT-J, a powerful generative model,

ensures that the results are reflective of state-of-the-art performance in language modeling tasks. The combination of DBpedia and GPT-J allows us to thoroughly investigate the behavior of the model across various ablation settings, ensuring robust insights into the performance of different methods and configurations.

B.1 Pruning for Classification

In our evaluation setup, we use pruning for classification by focusing only on the first token of the label, which is unique across all datasets. A common approach in the in-context learning setup is to iterate over all possible labels for each test sample and select the label with the highest probability according to the language model (LM). However, this approach can become computationally expensive, especially in cases where there are a large number of classes.

Similarly to (Ratner et al., 2022), and given that the first token in each dataset is unique, we predict only the first token of the label and perform classification based on this value. While this approach deviates slightly from the common practice of iterating over all possible labels, the effect on the results should be minor.

B.2 Test Set Size

For our experiments, we used a varying number of test examples depending on the dataset. Specifically, we used 100 test examples for the SST-2 dataset, and for datasets with a larger number of classes, the number of test examples was scaled linearly with the number of classes. For example, in the DBpedia dataset, which has 7 times more classes than SST-2, we used 700 test examples to ensure that the evaluation is proportional to the number of classes. This scaling helps to maintain a balanced evaluation across datasets with differing complexities, ensuring robust performance metrics for each method.

C Instruction Details

In some of the experiments, we use specific instructions to guide the model in performing the classification tasks. Below in table 4 that shows the instructions used for each dataset across all relevant methods:

D Dataset Details

In our experiments, we used four different datasets, each representing a unique classification task. Table 5 provides an overview of the datasets and their respective tasks. Each dataset has a varying number of classes, denoted by $|C|$, which are detailed below:

- **SST-2:** This dataset is used for *sentiment analysis*, where the task is to classify movie reviews as either positive or negative. It contains 2 distinct classes.
- **AG News:** The AG News dataset is used for *news classification*. The task is to classify news articles into one of four categories: World, Sports, Business, and Technology. This dataset contains 4 classes.
- **DBpedia:** The DBpedia dataset is focused on *ontology classification*. The task involves classifying textual content into one of 14 distinct categories, which include entities such as Company, Artist, Village, and more.
- **TREC:** This dataset is used for *question classification*, where the goal is to classify questions into one of 6 answer types, including Description, Entity, Human, and Location.

Each dataset contains a specific number of examples based on its classification task, allowing us to evaluate the model’s performance across a diverse range of challenges.

E Template Details

In our experiments, we use randomly selected templates from the options provided in table 6, suggested in (Voronov et al., 2024). Each dataset is associated with both input and output templates, which are used to format the input data and the expected output during few-shot learning tasks.

- **Input Template:** As shown, this column lists the different templates for formatting the input data. For example, the SST-2 dataset uses "input: " and "text: " as input templates to introduce the input text.
- **Intra-Separator:** This separator is used between components (input and output) within a single example. For instance, AG News uses "\n" as an intra-separator between the input sentence and the output label.
- **Output Template:** The output template defines how the expected output is structured. For example, SST-2 employs formats like "output: , target: , label: " to guide the model in generating structured output.
- **Inter-Separator:** This column represents the separator used between multiple examples during training. In datasets like AG News and DBpedia, "\n\n" is used to separate examples.

We randomly select templates from the ones listed in table 6 for each experiment. This randomness in selecting templates introduces variability in the prompts, making the evaluation more robust and testing the model’s ability to generalize across different input-output structures.

F Implementation Details

F.1 Hyperparameter Details

In table 7 we present the hyperparameters used in our experiments across different models and datasets. The table provides the specific learning rates ('lr'), epsilon values ('\epsilon'), and format settings for the various methods applied to each dataset. The experiments were conducted using multiple model architectures, including **BLOOM 1.7B**, **GPT-J 6B**, and **Llama3 8B**, and we selected

Dataset	Instruction
SST2	Classify the sentiment of the following text as positive or negative:
AG News	Classify the following text into one of the following categories: World, Sports, Business, Technology
DBpedia	Classify the following text into one of the following categories: Company, Educational Institution, Artist, Athlete, Office Holder, Mean Of Transportation, Building, Natural Place, Village, Animal, Plant, Album, Film, Written Work
TREC	Classify the following text into one of the following categories: Description, Entity, Expression, Human, Location, Number

Table 4: Instructions used for relevant datasets in the experiments.

the best hyperparameters for each experiment: 2, 4, and 6 shots. Below is an overview of the key hyperparameters:

- **Learning Rate ('lr'):** The table provides the learning rates used for each method and dataset combination. For methods like *Prefix Tuning (PT)*, *Prompt Tuning (PT)*, *IPT*, and *LoRA*, learning rates vary from **1e-5** to **1e-3**, depending on the specific model and dataset.
- **CPT Hyperparameters:** For *CPT*, we also report epsilon values (' ϵ ') for both the *input* and the *format* components. These epsilon values control the magnitude of the perturbations applied during optimization. The values of epsilon vary across different models and datasets, generally ranging from **1e-2** to **1e-0** for both input and format components.
- **Model Variability:** The table reflects variability in hyperparameter choices depending on the model size and architecture. For instance, *GPT-3 6B* typically requires higher learning rates compared to *BLOOM 1.7B*, as seen with *CPT* and other methods. The hyperparameters are carefully tuned to optimize performance on tasks such as SST-2, AG News, DBpedia, and TREC.

Dataset	Task	$ C $
SST-2	Sentiment analysis (movie)	2
AG News	News classification (topic)	4
DBpedia	Ontology classification	14
TREC	Question classification (answer type)	6

Table 5: **Dataset Overview** These are the datasets used, representing a range of different types of classification tasks, including SST-2, AG News, DBpedia, and TREC. Each dataset has a varying number of classes (denoted by $|C|$).

These hyperparameters are critical for achieving optimal performance in few-shot learning settings. They control the learning process, model updates, and how much the model is allowed to adapt to new data. The values in table 7 are based on extensive experimentation and fine-tuning to ensure the best results for each method and dataset.

F.2 Methods Implementation Details

In our experiments, we utilized existing implementations for several methods and implemented IPT ourselves. Specifically, we used the implementations provided by the *Parameter-Efficient Fine-Tuning (Mangrulkar et al., 2022) (PEFT)* library¹ for methods such as **LoRA**, **Prefix Tuning**, and **Prompt Tuning (PT)**. For IPT, we built our implementation based on the PEFT framework.

For all experiments, we used the recommended parameters:

- For LoRA, we set $\alpha = 16$ and the rank $r = 8$.
- For Prompt Tuning, Prefix Tuning, and IPT we used 8 learnable tokens.

By using the PEFT framework, we ensure that our fine-tuning processes for LoRA, Prefix Tuning, and PT are aligned with current standards, while our custom IPT implementation extends the framework to allow for additional flexibility in parameter-efficient training.

F.3 Training Details

We utilized the 'Fine-tune a pretrained model' package from (Wolf et al., 2020), which provides a comprehensive framework for training and evaluating models². For all baselines, we employed the default parameters provided by the trainer, ensuring

¹<https://huggingface.co/docs/peft/en/index>

²<https://huggingface.co/docs/transformers/en/training>

Dataset	Input Template	Intra-Separator	Output Template	Inter-Separator
SST-2	"input: {}", "text: ",	" "	"output: {}", "target: {}", "label: {}", "emotion: {}", "sentiment: {}", "A {} one.", "It was {}. ", "All in all {}. ", "A {} piece."	" "
AG News	"sentence: {}",	"\n"	"output: {}", "target: {}", "label: {}",	"\n"
DBpedia	"{}"		"Topic: {}. ", "Subject: {}. ",	"\n\n"
TREC			"This is about {}. ", "It is about {}. "	

Table 6: **Template Options for Various Datasets** We provide various template options for different datasets. Each dataset include both input and output templates, and also includes intra-separators between inputs and labels, as well as inter-separators between examples.

consistency across experiments. Each model was trained for 25 epochs, allowing sufficient time for convergence while maintaining uniform training conditions across methods.

G Input Preparation

In this section, we provide a detailed explanation of how the input is constructed for different methods, including Prompt Tuning (PT), Instruction Prompt Tuning (IPT), and Context-Aware Prompt Tuning (CPT), both with and without the \dagger variant. To clarify the differences, we use SST-2 as an example with the instruction: *"Classify the sentiment of the following text as positive or negative."*

Each example is constructed using a template that includes input: and output:, where the input corresponds to the actual text of the example, and the output corresponds to its label. For instance:

- **Example 1:** The input is "the greatest musicians", and the output is "positive".
- **Example 2:** The input is "the action is stilted", and the output is "negative".

Using the template, these examples are represented as:

- **Example 1:** input: the greatest musicians output: positive
- **Example 2:** input: the action is stilted output: negative

This template-based construction ensures consistency across the methods, allowing us to clearly define how the input and output are represented in different approaches, such as PT, IPT, and CPT.

table 8 outlines the construction of the prefix for each method and highlights which parts are updated during training.

H Projected Gradient Descent (PGD) Algorithm

In our method, we initialize the context tokens, denoted as x_i , using the training examples, with each token x_i associated with a vector δ_i , which is initially set to zero. For simplicity, we use x_i and δ_i to denote these components only in this part of the explanation.

During the optimization process, the tokens x_i remain fixed, while the δ_i vectors are updated iteratively. After each optimizer update, we perform a post-processing step where each δ_i is projected to ensure that its L2 norm does not exceed a pre-defined limit, ϵ . It is important to note that this projection step is independent of the optimizer and serves as an additional operation to control the extent of change for each context token.

- 1: Initialize each $\delta_i \leftarrow 0$
- 2: Initialize $x_i \leftarrow \text{training_examples_tokens}$
- 3: **for** $j \leftarrow 1$ to $\text{num_of_training_steps}$ **do**
- 4: $\delta_i \leftarrow \delta_i - \alpha \nabla \text{Loss}(f(x_i + \delta_i), y_i)$ \triangleright Gradient descent step
- 5: $n_i \leftarrow \|\delta_i\|$ \triangleright Compute the L2 norm of δ_i
- 6: $\delta_i \leftarrow \delta_i \times \text{clip}(n_i, \epsilon) / n_i$ \triangleright Project δ_i to ensure L2 norm $\leq \epsilon$
- 7: **end for**

This ensures that the updates to δ_i remain constrained, preventing excessive modifications to the context tokens and maintaining a balance between optimization and regularization. The process allows the model to adapt while ensuring that changes to the context tokens remain meaningful and controlled.

I Evaluating the Impact of Projected Gradient Descent (PGD)

Our method use the same optimizer used for all baselines. However, our method incorporates an additional step after each parameter update: we

Dataset	Method	Parameter	BLOOM 1.7B			GPT-J 6B			Llama3 8B		
			2	4	6	2	4	6	2	4	6
SST-2	Prefix Tuning	lr	$1e-3$	$1e-3$	$1e-3$	$1e-5$	$1e-4$	$1e-3$	—	—	—
	PT†	lr	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-3$	$1e-3$	$1e-5$	$1e-5$	$1e-5$
	PT	lr	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$
	IPT†	lr	$1e-5$	$1e-4$	$1e-4$	$1e-5$	$1e-3$	$1e-4$	$1e-5$	$1e-5$	$1e-4$
	IPT	lr	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
	LoRA	lr	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-4$
	CPT†	lr	$1e-5$	$1e-3$	$1e-4$	$1e-5$	$1e-4$	$1e-3$	$1e-5$	$1e-5$	$1e-5$
		Input ϵ	$1e-3$	$1e-0$	$1e-0$	$1e-3$	$1e-1$	$1e-1$	$1e-1$	$1e-1$	$1e-0$
		Format ϵ	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-2$	$1e-3$	$1e-2$	$1e-1$	$1e-0$
	CPT	lr	$1e-3$	$1e-3$	$1e-4$	$1e-5$	$1e-4$	$1e-4$	$1e-3$	$1e-4$	$1e-4$
		Input ϵ	$1e-2$	$1e-0$	$1e-0$	$1e-3$	$1e-0$	$1e-0$	$1e-2$	$1e-0$	$1e-2$
		Format ϵ	$1e-2$	$1e-2$	$1e-3$	$1e-3$	$1e-3$	$1e-2$	$1e-3$	$1e-3$	$1e-3$
AG News	Prefix Tuning	lr	$1e-4$	$1e-3$	$1e-3$	$1e-5$	$1e-5$	$1e-3$	—	—	—
	PT†	lr	$1e-3$	$1e-3$	$1e-3$	$1e-5$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$
	PT	lr	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-3$	$1e-3$	$1e-4$	$1e-5$	$1e-4$
	IPT†	lr	$1e-3$	$1e-3$	$1e-3$	$1e-5$	$1e-4$	$1e-4$	$1e-4$	$1e-5$	$1e-5$
	IPT	lr	$1e-4$	$1e-3$	$1e-4$	$1e-5$	$1e-5$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
	LoRA	lr	$1e-5$	$1e-4$	$1e-3$	$1e-5$	$1e-5$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
	CPT†	lr	$1e-4$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
		Input ϵ	$1e-2$	$1e-0$	$1e-0$	$1e-1$	$1e-1$	$1e-2$	$1e-1$	$1e-3$	$1e-3$
		Format ϵ	$1e-1$	$1e-2$	$1e-0$	$1e-1$	$1e-3$	$1e-0$	$1e-1$	$1e-2$	$1e-3$
	CPT	lr	$1e-4$	$1e-4$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-3$	$1e-4$	$1e-3$
		Input ϵ	$1e-2$	$1e-0$	$1e-0$	$1e-2$	$1e-0$	$1e-0$	$1e-2$	$1e-3$	$1e-3$
		Format ϵ	$1e-2$	$1e-0$	$1e-0$	$1e-3$	$1e-3$	$1e-0$	$1e-3$	$1e-3$	$1e-3$
DBpedia	Prefix Tuning	lr	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-3$	—	—	—
	PT†	lr	$1e-3$	$1e-5$	$1e-3$	$1e-5$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$
	PT	lr	$1e-4$	$1e-5$	$1e-4$	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-5$	$1e-5$
	IPT†	lr	$1e-4$	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-5$	$1e-5$	$1e-5$	$1e-5$
	IPT	lr	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$
	LoRA	lr	$1e-4$	$1e-5$	$1e-5$	$1e-4$	$1e-4$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
	CPT†	lr	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-5$
		Input ϵ	$1e-2$	$1e-2$	$1e-1$	$1e-0$	$1e-1$	$1e-1$	$1e-0$	$1e-1$	$1e-1$
		Format ϵ	$1e-1$	$1e-0$	$1e-1$	$1e-3$	$1e-0$	$1e-1$	$1e-1$	$1e-0$	$1e-1$
	CPT	lr	$1e-4$	$1e-4$	$1e-5$	$1e-4$	$1e-4$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
		Input ϵ	$1e-0$	$1e-2$	$1e-0$	$1e-0$	$1e-0$	$1e-0$	$1e-2$	$1e-0$	$1e-3$
		Format ϵ	$1e-0$	$1e-0$	$1e-0$	$1e-0$	$1e-3$	$1e-3$	$1e-2$	$1e-3$	$1e-2$
TREC	Prefix Tuning	lr	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-5$	—	—	—
	PT†	lr	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$
	PT	lr	$1e-5$	$1e-3$	$1e-3$	$1e-5$	$1e-3$	$1e-3$	$1e-5$	$1e-5$	$1e-5$
	IPT†	lr	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-3$	$1e-4$	$1e-4$	$1e-4$	$1e-5$
	IPT	lr	$1e-5$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$	$1e-5$	$1e-5$	$1e-5$
	LoRA	lr	$1e-4$	$1e-5$	$1e-5$	$1e-5$	$1e-5$	$1e-4$	$1e-5$	$1e-5$	$1e-4$
	CPT†	lr	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-x$	$1e-4$	$1e-5$	$1e-5$
		Input ϵ	$1e-0$	$1e-0$	$1e-0$	$1e-1$	$1e-0$	$1e-0$	$1e-1$	$1e-1$	$1e-1$
		Format ϵ	$1e-3$	$1e-1$	$1e-2$	$1e-1$	$1e-0$	$1e-2$	$1e-3$	$1e-0$	$1e-0$
	CPT	lr	$1e-3$	$1e-3$	$1e-4$	$1e-3$	$1e-3$	$1e-3$	$1e-4$	$1e-4$	$1e-4$
		Input ϵ	$1e-0$	$1e-0$	$1e-0$	$1e-0$	$1e-0$	$1e-3$	$1e-0$	$1e-0$	$1e-0$
		Format ϵ	$1e-0$	$1e-0$	$1e-3$	$1e-2$	$1e-2$	$1e-0$	$1e-2$	$1e-3$	$1e-0$

Table 7: **Hyperparameters** Hyperparameters used for each experiment across 2, 4, and 6 shots for different models, including BLOOM 1.7B, GPT-J 6B, and Llama3 8B. The table shows learning rates (lr), epsilon values for input and format, and other parameters for methods such as Prefix Tuning, Prompt Tuning, IPT, LoRA, and CPT. The experiments were conducted on datasets like SST-2, AG News, DBpedia, and TREC.

project each token, restricting its allowed change. The allowed change is determined by the hyperparameters Input ϵ and Format ϵ , which define the L2 norm limit for each token’s modification.

To ensure that PGD (Madry et al., 2017) is not the sole reason for our method’s improvement, we

conducted two types of experiments. First, we compared our method without PGD to PT and IPT. Second, we added a PGD step to PT and IPT for comparison.

For the first experiment, we compared CPT (without PGD) to PT and IPT on the DBpedia

Method	Prefix Construction
PT	In this part, we use only random embedding initialization.
PT [†]	Classify the sentiment of the following text as positive or negative.
IPT	In this part, we use only random embedding initialization. input: the greatest musicians output: positive. input: the action is stilted output: negative.
IPT [†]	Classify the sentiment of the following text as positive or negative. input: the greatest musicians output: positive. input: the action is stilted output: negative.
CPT	input: the greatest musicians output: positive. input: the action is stilted output: negative.
CPT [†]	Classify the sentiment of the following text as positive or negative. input: the greatest musicians output: positive. input: the action is stilted output: negative.

Table 8: Input Construction for PT, IPT, and CPT (with and without [†]) using SST-2. The updated text during training is marked in red.

dataset. The results for 2, 4, and 6 shots are presented in Table 9.

Method	2 Shots	4 Shots	6 Shots
PT	23.39	29.69	40.53
IPT	52.86	67.27	70.73
CPT (No PGD)	68.28	74.17	77.52

Table 9: Performance Comparison Without PGD (DBpedia), using GPT-J.

For the second experiment, we compared CPT[†] to PT[†] and IPT[†] (with and without PGD) on the DBpedia dataset. To ensure a fair comparison, we performed hyperparameter tuning (HPT) over ϵ and the learning rate for both PT and IPT. The results for 2, 4, and 6 shots are presented in Table 10.

Method	2 Shots	4 Shots	6 Shots
PT [†]	12.96	22.12	37.44
PT [†] + PGD	12.80	22.02	38.69
IPT [†]	47.10	66.37	75.09
IPT [†] + PGD	47.10	66.40	75.09
CPT [†] + PGD	52.87	77.30	81.00

Table 10: Performance Comparison With and Without PGD (DBpedia), using GPT-J.

The results clearly demonstrate that, in both experiments, our method consistently outperforms PT and IPT. Furthermore, it is evident that other methods do not necessarily benefit from the addition of PGD. While we cannot definitively explain this, we hypothesize that it may be due to the highly ef-

fective way in which we employ PGD, leveraging prior knowledge about the structure of the input, format, and labels within the context. Our approach allows us to apply distinct projections to different components of the context, which we believe significantly contributes to the superior performance of our method.