Large Language Model Prompting With Episodic Memory

Dai Do[a,*]**, Quan Tran**[b]**, Svetha Venkatesh**[a] **and Hung Le**[a]

[a]Applied AI Institute, Deakin University, Australia
[b]ServiceNow Research, USA

Abstract.   Prompt optimization is essential for enhancing the performance of Large Language Models (LLMs) in a range of Natural Language Processing (NLP) tasks, particularly in scenarios of few-shot learning where training examples are incorporated directly into the prompt. Despite the growing interest in optimizing prompts with few-shot examples, existing methods for prompt optimization are often resource-intensive or perform inadequately. In this work, we propose PrOmpting with Episodic Memory (POEM), a novel prompt optimization technique that is simple, efficient, and demonstrates strong generalization capabilities. We approach prompt optimization as a Reinforcement Learning (RL) challenge, using episodic memory to archive combinations of input data, permutations of few-shot examples, and the rewards observed during training. In the testing phase, we optimize the sequence of examples for each test query by selecting the sequence that yields the highest total rewards from the top-k most similar training examples in the episodic memory. Our results show that POEM outperforms recent techniques like TEMPERA and RLPrompt by over 5.3% in various text classification tasks. Furthermore, our approach adapts well to broader language understanding tasks, consistently outperforming conventional heuristic methods for ordering examples.

1   Introduction

The recent rapid advancements in Language Models (LMs) have underscored the increasing significance of utilizing pre-trained language models, especially when paired with appropriate prompts, as evidenced by seminal works [4, 1, 5, 9]. As Language Models increase in parameter count, they unveil new capabilities such as In-Context Learning (ICL) [4], enabling Large Language Models (LLMs) to tackle tasks with just a few example demonstrations in the prompt. ICL offers a data-efficient approach for performing Natural Language Processing (NLP) tasks, achieving remarkable few-shot performances across many downstream tasks [20, 10, 22, 23].

However, the prompt content and ICL examples necessitate meticulous tuning to ensure consistent performance across various tasks. To optimize prompt contents, early attempts focus on tuning the embeddings via gradient descent ("soft prompts", [18, 23]). Unfortunately, soft prompts require gradients from LLMs to construct prompts and often face challenges with interpretability and quality [13, 18]. Additionally, they struggle to handle ICL examples within the prompt, and thus can only be used for zero-shot prompting. Consequently, the current state-of-the-art has shifted towards discrete prompt optimization [7, 44], which enhances interpretability and permits ICL optimization.

Selecting the right in-context examples and their orders in prompts is crucial for ICL optimization [22]. This task is challenging due to the vast array of possible combinations and diverse instructions [26]. Furthermore, the arrangement of these examples may introduce biases, including majority, recency, and primacy biases [26, 46, 28]. Although reasonable example selection can be achieved through nearest neighbor retrieval [22], determining the optimal order of examples remains an unresolved research challenge.

Initial efforts employed heuristic rules to rank examples in descending or ascending order based on their similarity to the test instance [22, 26]. More recent work attempts to use LLMs as black-box optimizer [37, 29], calibration [46] or applies Reinforcement Learning (RL)-based method for generating prompts with ICL examples [44].

Discrete prompt optimization presents its own set of challenges. Heuristic methods lack optimization principles, leading to success in some cases but failure in others [22]. Black-box methods are guided-optimization and gradient-free. However, they are query-agnostic, thus failing to incorporate any query-related context into the prompt. This leads to downstream performance degradation [41]. Moreover, they often require additional LLM computation for prompt generation, resulting in extensive resource usage [42, 30, 47]. Although using RL-based methods for prompt editing sequentially presents a potential solution that does not require extra LLM for prompt generation [44], their slow convergence and intrinsic complexity hinder effectiveness.

In this paper, we propose a novel and efficient memory-based approach to optimize the order of ICL examples within LLM prompts. Drawing inspiration from the rapid, model-free, and instance-based learning evident in the hippocampus region of the human brain [17], our method eliminates the necessity for complex reinforcement learning optimization while being more reliable than heuristic methods through performance-driven optimization. Leveraging episodic control mechanisms in reinforcement learning [3, 14, 16], we formulate each evaluation of training data as an episode and utilize an episodic memory to store the performance of any combination of the training data and ICL orderings. By sampling and evaluating certain training inputs and ICL order pairs, we avoid exhaustive searches across all data-ICL order combinations, which is particularly beneficial when LLM evaluation is costly. During testing, this memory serves as a non-parametric nearest-neighbors model, utilizing the recorded performance of similar training data to determine the optimal order for the testing data.

To ensure the robust generalization of our episodic memory, we devise specialized representations for both the text input and the ordering of ICL examples. Specifically, we encode the input using the last hidden states of a pre-trained language model, ensuring high-quality similarity-based retrieval during testing. Moreover, directly
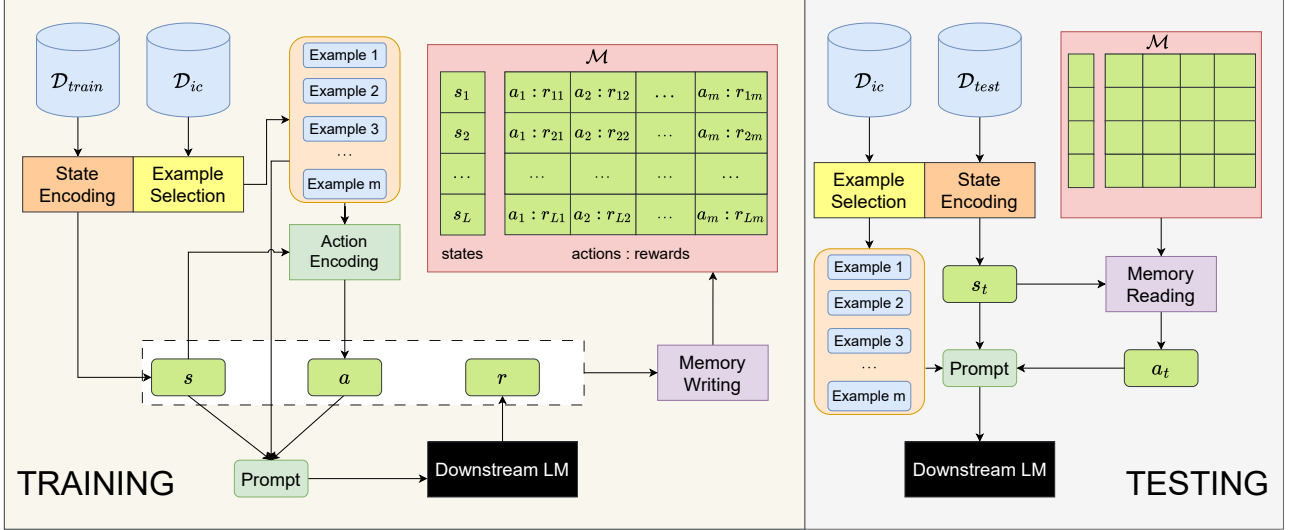
Figure 1. POEM Architecture. Training (left): In this phase, we select examples from the in-context dataset D_{ic}. The training query and the ICL example ordering are encoded into s and a, respectively, and are used to construct a prompt for each training query. Then, we receive a reward r by feeding the prompt to the downstream language model (LM), and we store the state, action, and reward in memory $\mathcal{M}$ using **Memory Writing** (Eq. 9). Testing (right): During this phase, for each testing query s_t, we conduct **Memory Reading** using nearest neighbor estimation to get the action with the highest estimated value (Eq. 10). We then build the prompt for the test query by producing the ICL examples that correspond to the best ordering action a_t.

encoding the permutation as a sequence of ICL examples would result in a vast search space. Suppose there are M training samples in total and we can select m examples for each prompt, we would have $\frac{M!}{(M-m)!}$ possible ICL arrangements. Instead, we represent the ordering as a permutation of the similarity rank of the in-context examples. Specifically, given a specific ordering of ICL examples, we (1) gauge the similarity between the examples to the testing input, (2) assign a rank to each example based on the similarity, and (3) encode the ICL ordering as the sequence of the ranks. Our encoding operates within the similarity rank space rather than the text space, effectively reducing the search space from $\frac{M!}{(M-m)!}$ to $m!$. More importantly, this approach also encourages generalization as the permutation focuses solely on the arrangement of the example rank rather than specific content.

In summary, we propose POEM, a method that optimizes in-context example ordering during test time using an Episodic Memory. It utilizes a similarity ranking to encode example orders based on their proximity to the test instance. In our few-shot classification experiments across seven datasets, POEM outperforms TEMPERA on six datasets and surpasses RLPrompt on all, with a significant performance margin of 13.4%. For tasks requiring larger LLMs, such as Commonsense Reasoning and Question Answering, POEM consistently outperforms heuristic baselines across all four LLMs tested.

2 Method

2.1 Problem Formulation

Few-shot text classification. We follow the standard few-shot setting for downstream tasks of language models [4]. In this setup, a pretrained language model $\mathcal{L}$ is paired with a task-specific dataset $\mathcal{D}$, where $\mathcal{Y}$ represents the label space. $\mathcal{Y}$ may vary in form; it can be categorical, as in classification tasks, or sequential, as in question answering and commonsense reasoning tasks. For classification tasks, we randomly assemble a dataset of $L = g * \mathcal{G}$ samples, with g as the samples per label and $\mathcal{G}$ as the total labels. In cases without categori-

cal labels, we randomly sample L instances from $\mathcal{D}$. This forms the training dataset, denoted as $D_{train} = \{x_i, y_i\}_{i=1}^{L}$, while a separate hold-out set D_{test} is reserved for evaluation.

In-context Learning. Following GPT-3 paper [4], *In-context Learning is a paradigm that allows language models to learn tasks given only a few examples in the form of demonstration.* Given a test sentence x_{test}, template-based construction Ψ, along with an optional task description t_{desc} and a set of in-context examples $\mathcal{T} = \{x_i, y_i\}_{i=1}^{m}$, denoting Γ as the prompt construction function, we can formulate an input prompt p as follows:

$$
\begin{aligned}
p &= \Gamma\left(t_{desc}, \mathcal{T}, x_{test}\right) \\
&= t_{desc} \oplus \Psi\left(x_1, y_1\right) \oplus \cdots \oplus \Psi\left(x_m, y_m\right) \oplus \Psi\left(x_{test}, *\right)
\end{aligned}
\tag{1}
$$

where $\oplus$ is the concatenate function and m is the number of in-context examples for each prompt. In line with the few-shot setting [44, 7], we consider an in-context set of M samples, denoted as $\mathcal{D}_{ic} \in \mathcal{D}$, excluding the few-shot training data. The in-context examples will be sampled only from this set.

In-context learning allows the construction of the task's output distribution $p_{LM}(y|x, p)$, where x and y are string input and output, respectively. This approach enables in-context learning to actively influence the selection and arrangement of the demonstration set $\mathcal{T}$ within the input prompt p, thereby refining its capacity for effective optimization.

Reinforcement Learning Formulation. We formulate in-context prompt optimization as an RL problem where a state $\mathbf{s} = \mathcal{E}(x)$ is the embedding of the input x, where $\mathcal{E}$ is a pre-trained encoder. During training, given a set of m in-context examples, the RL agent selects one of the possible permutations as its action a from the action space $\mathcal{A}$. We then construct the prompt p using the default task description/instruction (if any), combined with the a-ordered in-context examples, and query it to the downstream LM to get the reward r. The goal of the agent is to learn a policy that maximizes the episodic return $R_t = \sum_{h=0}^{H-t} r_{t+h}$, where H is the time step at

Algorithm 1 In-context examples order optimization with Episodic Memory

---

Require: Language model $\mathcal{L}$, State encoder $\mathcal{E}$, Training set $\mathcal{D}_{train}$, Evaluation set $\mathcal{D}_{test}$, In-context set $\mathcal{D}_{ic}$, Number of iterations N, Episodic Memory $\mathcal{M}$, Number of neighbors k, Task description t_{desc}, Template transformation Ψ, Number of in-context examples per prompt m, Prompt construction function Γ

1: Initialize $\mathcal{M} = \emptyset$
2: **Training:**
3: **for** episode $n = 1$ to N **do**
4:     Random sample batch $\mathcal{B} \sim D_{train}$
5:     **for** x $\in \mathcal{B}$ **do**
6:         Receive state $s = \mathcal{E}(x)$
7:         Receive in-context examples $\mathcal{T}_s = \Omega(s, \mathcal{D}_{ic})$
8:         Select permutation $a \leftarrow \epsilon$-greedy policy using $\mathcal{M}$ (Eq. 10)
9:         Reorder the in-context examples set $\mathcal{T}_s^a = \nabla(\mathcal{T}_s, a)$
10:        Get prompt $p = \Gamma(t_{desc}, \Psi(\mathcal{T}_s^a))$
11:        Receive reward r
12:        Update $\mathcal{M}$ using **Memory writing** with r (Eq. 9)
13:     **endfor**
14: **endfor**
15: **Testing:**
16: **for** $x_t \in \mathcal{D}_{test}$ **do**
17:     Receive state $s_t = \mathcal{E}(x_t)$
18:     Receive in-context example set $\mathcal{T}_{s_t} = \Omega(s_t, \mathcal{D}_{ic})$
19:     Obtain $\widehat{\mathcal{M}}(s_t, a)$ with **Memory reading** (Eq. 10)
20:     Obtain a_t via Eq. 11
21:     Reorder the in-context examples set $\mathcal{T}_{s_t}^{a_t} = \nabla(\mathcal{T}_{s_t}, a_t)$
22:     Get prompt $p = \Gamma(t_{desc}, \Psi(\mathcal{T}_{s_t}^{a_t}))$
23:     Get prediction $\hat{y} = \mathcal{L}(x_t, p)$
24: **endfor**

---

which the episode ends. To simplify the formulation, our episode consists of only one step, during which the sole action is selecting the permutation of the in-context examples.

2.2 Prompting with Episodic Memory

In this section, we present the architecture of our episodic memory. The memory is structured as a dictionary, storing the embeddings of training sentences (states) as keys. Each key's value is another dictionary, mapping a permutation (action) to its respective reward. We denote a key by s_i, and for each key, a_j represents the j-th permutation, while r_{ij} denotes the associated reward. Our episodic memory $\mathcal{M}$ can be represented by the following structure:

$$\mathcal{M} = \{s_i : \{a_1 : r_{i1}, a_2 : r_{i2}, \ldots, a_p : r_{ip}\}\}_{i=1}^{L} \quad (2)$$

Here, p signifies the total number of permutations available for the in-context examples (with m examples per prompt, $p = m!$). L is the total of the stored states in the episodic memory. An illustration of our memory architecture is given in Fig. 1. Below are the detailed components of the memory.

State representation Obtaining text representations for the state is crucial for memory storage and retrieval. In our approach, we employ the encoder $\mathcal{E}$ as SentenceTransformers, which is designed to generate high-quality sentence embeddings that can be used effectively in various NLP tasks [31].

Example selection To simplify the optimization, we do not aim to optimize the example selection process. Therefore, following prior work [22], from the in-context dataset D_{ic}, given input x, we select

m in-context examples that are semantically closest to x. We measure the semantic similarity between in-context example x_{ic}^i and the input query x using Cosine Similarity:

$$CS\left(x, x_{ic}^i\right) = \frac{s \cdot s_{ic}^i}{||s||||s_{ic}^i||} \quad (3)$$

In the context of a *few-shot classification* task, where label biases can significantly influence the prompting outcome, it is crucial to maintain an equal number of examples for each label. To address this, we propose the following strategy: if there are $\mathcal{G}$ unique labels in $\mathcal{D}$, and $\mathcal{G} < m$, we select $\frac{m}{\mathcal{G}}$ closest samples from each label. Conversely, if $\mathcal{G} \geq m$, we iteratively choose one sample for each label until we have sufficient in-context examples. We denote the above process as the function Ω that retrieves context examples for each state s:

$$\mathcal{T}_s = \Omega(s, \mathcal{D}_{ic}) \quad (4)$$

where $\mathcal{T}_s$ is the set of in-context example for the state s

Action encoding An action refers to a specific arrangement of in-context examples within $\mathcal{T}_s$. As mentioned in the introduction, a naive approach of encoding the action as a sequence of ICL examples will lead to a huge action space. Therefore, we propose to encode the action as a sequence of similarity ranks. Concretely, for each state s, we measure the semantic similarity between it and the states of the in-context examples using Eq. 3. Next, we rank each example according to its similarity to the input query and encode the ICL ordering as a sequence of these ranks. For example, consider an action represented as $a = (1, m, \ldots, 3)$. This permutation indicates that the first in-context example is the closest to the test query, the second is the farthest, and the sequence continues until the m-th position, which is the third closest example. This approach encapsulates the relational similarity information among in-context examples and the test query, with the action space size being $m!$.

Given a set of in-context examples $\mathcal{T}_s$ and an action a, we define a permutation function ∇, used to reorder the elements of $\mathcal{T}_s$ according to the sequence specified by a. The action a represents a specific ordering of indices that correspond to the elements in $\mathcal{T}_s$. Formally, we have:

$$\mathcal{T}_s^a = \nabla(\mathcal{T}_s, a) \quad (5)$$

where $\mathcal{T}_s^a$ represents the sequence of in-context examples rearranged in accordance with action a.

Reward design For *classification* tasks that use Masked LM such as RoBERTa [25], for each query x, we define the reward based on the log probability of the output label of the model $\log P_{\mathcal{L}}(\hat{y}|x, p)$

$$r(c, x, p) = \lambda_1 \log \mathcal{P}_{\mathcal{L}}(\hat{y}_c \mid x, p) - \lambda_2 \arg\max_{c \neq c'} \log \mathcal{P}_{\mathcal{L}}(\hat{y}_{c'} \mid x, p) \quad (6)$$

where $\lambda_1 > 0$ and $\lambda_2 > 0$ are the two balance hyperparameters for the positive and negative terms respectively.

For *generative* tasks (e.g: Commonsense Reasoning, Question Answering) that use Causal LM such as Llama [38], where ground truth label c is a sequence of tokens, we define the reward as:

$$r(c, x, p) = \lambda_1 \sum_{i=1}^{|c|} \log \mathcal{P}_{\mathcal{L}}(i \mid x, p) -$$
$$\lambda_2 \sum_{j=1}^{|c'|} \arg\max_{c \neq c'} \log \mathcal{P}_{\mathcal{L}}(j \mid x, p) \quad (7)$$

In the equations above, c' is the token/sequence that has the highest probability, other than the ground truth label c. In Eq. (7), i and j

are the tokens of two sequences c and c', respectively. Intuitively, for classification tasks, the reward is positive when the prediction is correct and negative otherwise.

For tasks that use Causal LM and Exact Match as evaluation metrics, we define the reward as:

$$r\left(c, x, p\right) = \begin{cases} 1, & \text{if } c = c' \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

2.3 Episodic Memory Operation

Our memory operation involves two main phases. **Training**: The RL agent interacts with the LLM to try actions and collect rewards for (s, a) pairs to fill the memory. **Testing**: The RL agent uses the memory to determine the ICL ordering for testing data.

Training In this phase, we collect and store the rewards defined above for state action pairs (s, a). Given an input query $x \in \mathcal{D}_{train}$, we get its representation using the pre-trained encoder $s = \mathcal{E}(x)$. Given the a pair of state and action (s, a) and the reward r, following [3], our Episodic Memory $\mathcal{M}$ is updated using **Memory writing** as follows:

$$\mathcal{M}(s, a) \leftarrow \begin{cases} r, & \text{if } (s, a) \notin \mathcal{M} \\ \max\left\{\mathcal{M}(s, a), r\right\}, & \text{otherwise} \end{cases} \quad (9)$$

It is noted that the stored rewards never decrease, indicating a focus on high-return actions. In our setting, we use the highest possible reward that the downstream LM can achieve for the pair (s, a) to estimate the value of using ICL ordering a for input x. The motivation behind this is to emulate the brain's specialized learning mechanisms that exploit predictable patterns in the environment, enabling rapid learning from high-return actions recorded in memory.

We limit the size of our memory to be equal to the size of the training dataset $\mathcal{D}_{train}$. In the few-shot scenario where there are limited training samples, it is guaranteed that our memory will not be overflowed. However, in other settings where $\mathcal{D}_{train}$ is too big, our memory can have a fixed smaller size, and when a new state-value pair has to be introduced, the least recently used state will be discarded.

Testing In this phase, we select the optimal permutation for each testing query. In reality, it is common to receive novel states (states that are not seen in training). To handle this, we employ a nearest-neighbor estimator. In particular, we obtain the approximated value of a testing state s_t and a candidate action a using the following **Memory Reading** process [15]:

$$\widehat{\mathcal{M}}(s_t, a) = \begin{cases} \sum_{i=1}^{k} \mathcal{M}(s^i, a) \frac{CS(x_t, x^i)}{\sum_{j=1}^{k} CS(x_t, x^j)}, & \text{if } (s_t, a) \notin \mathcal{M} \\ \mathcal{M}(s_t, a), & \text{otherwise} \end{cases} \quad (10)$$

where s^i, $i = 1, \dots, k$ are the k states with the highest similarity to the testing state s_t, $CS(x_t, x^i)$ is the Cosine Similarity between the neighbor x^i to the test query x_t. The motivation behind the weighted sum is that we believe the semantic similarity between training data and the test query should correspondingly affect its weight in the ordering process. After having calculated $\widehat{\mathcal{M}}(s_t, a)$, we get the action a_t for s_t as follow:

$$a_t = \arg\max_a \widehat{\mathcal{M}}(s_t, a) \quad (11)$$

We note that this nearest neighbor estimation is a different process from the nearest neighbor retrieval in the example selection described in Section 2.2. Algo. 1 summarizes the procedure of our method.

3 Experiments

3.1 Few-Shot Text Classification

Datasets. For classification tasks, we conduct experiments for sentiment analysis, Natural Language Understanding (NLU), topic classification, and natural language inference datasets. For sentiment analysis, we choose SST-2 [35], IMDB [27] and CR [11] as datasets. For NLU task, we select COLA [39]. For Reading Comprehension, we choose Boolq [6]. For topic classification, we use AG News [45]. For Natural Language Inference (NLI), we choose QNLI [40]. The statistics, manual templates, and label words of this dataset are shown in Appendix A.3.

Baselines. For few-shot classification tasks, we compare our work with previous continuous and discrete prompt optimization methods. *Finetuning* is the method that finetunes the entire language model with a classification head using the few-shot dataset. For *Manual Prompt*, we use the hand-written prompts from [34]. *In-context Demonstration* randomly selects one example and concatenates it with the test query. *Black-box Tuning* [37] combines a discrete component optimized through non-gradient methods and a soft component refined via gradient descent. *RLPrompt* [7] generates discrete prompt tokens using RL framework. *TEMPERA* [44] trains a RL agent to edit prompt sequentially. To ensure a fair comparison, we rerun these methods under our setting. Besides complicated baselines, we also investigate a simple heuristic baseline. *Random Ordering* adopts nearest neighbor example selection, randomly permutes the in-context examples, and concatenates with the test query.

Experiment Setup. We use RoBERTa-large [25] as the downstream LM. To ensure a fair comparison, we follow the same setting from [7, 44] by testing all baselines on a few-shot text classification setting. Training dataset $\mathcal{D}_{train}$ has 16 examples per class, and we sample another 16 data points as in-context dataset $\mathcal{D}_{ic}$ ($M = 16$). For reporting the testing results, we make use of the models having the highest performance on the validation set and do inference on the test set $\mathcal{D}_{test}$ provided by the task. We use $m = 4$ examples for each prompt as in prior works [44]. For each run, we use the same training and in-context samples for all baselines. At test time, we select the number of nearest-neighbor $k = 10$ for POEM. We describe more details in terms of the instructions in Tab. 9 and templates in Tab. 10.

Results. We present our few-shot text classification results over 4 runs in Tab. 1. We can see that on most tasks, POEM outperforms previous baselines by a large margin. For example, we have a 2.9 % absolute gain on SST-2 task (over TEMPERA), and 5.6 % on IMDB, and the performance is almost comparable to finetuning the LM on QNLI task. Especially, for NLI and NLU tasks, on harder datasets like Boolq and COLA, we have a significance gain of 11.2 % and 14.2 % respectively. We also see that POEM has a much smaller variance between different runs than all other baselines, which indicates that it is more stable across different datasets. Unlike TEMPERA, which requires sequential modification of in-context examples, our method, POEM, enables one-step optimization. Our approach not only simplifies the process but also overcomes TEMPERA's slow inference involving several numbers of edits for each prompt. POEM also outperforms Random Ordering significantly, by nearly 13% on average, highlighting the importance of ICL optimization. More comparisons with advanced heuristic approaches are given in Appendix A.1.

3.2 General Language Understanding Tasks

We further extend experiments of our method to general language understanding tasks that require stronger downstream LLM to generate

	SST-2	Boolq	CR	AG News	IMDB	QNLI	COLA	**Average**
Finetune	80.6(3.9)	55.3(3.1)	73.3(7.5)	**84.9(3.6)**	85.4(5.2)	<u>55.4(1.6)</u>	55.6(9.9)	70.1
Manual Prompt	82.8	61.0	79.6	76.9	85.0	51.0	32.0	66.9
In-Context Demo.	85.8(0.7)	58.3(0.4)	85.5(1.5)	74.9(0.8)	69.8(0.6)	53.5(0.5)	56.0(1.2)	69.1
Instruction	89.0	61.0	80.8	54.8	<u>89.0</u>	**56.0**	<u>61.0</u>	70.2
Black-box Tuning	89.1(0.9)	53.3(1.7)	87.4(1.0)	83.5(0.9)	68.9(5.1)	50.8(0.6)	51.4(2.1)	69.2
RLPrompt	87.7(3.6)	41.6(3.0)	90.4(1.7)	73.8(5.8)	57.3(8.9)	50.1(0.8)	51.6(1.5)	64.6
TEMPERA	<u>90.5(1.6)</u>	54.9(5.8)	<u>91.1(1.1)</u>	81.6(2.5)	85.3(2.8)	51.6(1.0)	54.2(5.5)	<u>72.7</u>
Random Ordering	81.5(0.9)	<u>62.7(0.9)</u>	87.2(0.4)	57(0.2)	74.1(0.5)	53.9(0.4)	41.63(1.3)	65.4
POEM (Ours)	**93.4(0.2)**	**66.1(0.1)**	**92.6(0.2)**	80.3(1.1)	**90.9(0.1)**	54.3(0.2)	**68.4(0.4)**	**78**

Table 1. Accuracy and standard derivation (if available) of different baselines on few-shot classification over 4 seeds. Some methods like Manual Prompt produce the same results across seeds, and thus, have no standard derivation to report. The highest accuracy is **bolded** and the second-highest accuracy is <u>underlined</u>. The last column shows the average accuracy across all datasets in this table.

Model name	Baseline	Hellaswag	PIQA	TruthfulQA	TriviaQA
		acc.	acc.	acc.	EM
	0-shot	57.79	76.38	25.99	18.90
	Random	58.59	77.37	32.38	58.65
	Ascending	58.36	76.22	26.94	57.06
	Descending	<u>59.61</u>	<u>78.73</u>	<u>38.90</u>	<u>59.01</u>
Llama 2 7B	POEM (Ours)	**59.65**	**79.27**	**38.91**	**59.06**
	0-shot	60.73	77.58	24.63	27.26
	Random	62.40	<u>79.05</u>	32.79	65.98
	Ascending	61.39	76.77	26.94	64
	Descending	**63.17**	79	<u>40.41</u>	**67.13**
Llama 2 13B	POEM (Ours)	<u>63.13</u>	**79.43**	**40.95**	<u>67.06</u>
	0-shot	63.76	79.27	28.16	42.69
	Random	65.59	80.72	38.64	73.09
	Ascending	65.59	80.79	32.24	72.07
	Descending	<u>66.99</u>	**81.66**	<u>45.17</u>	<u>73.53</u>
Llama 2 70B	POEM(Ours)	**67.15**	<u>81.12</u>	**45.71**	**73.56**

Table 2. Comparison between POEM and other heuristic baselines. All baselines (except for 0-shot) use 4 in-context examples for each prompt. **Bolded** are the best results and <u>underlined</u> are the second best. The metric for evaluation for each dataset is written below the dataset names. *acc.* stands for accuracy and *EM* stands for exact match.

Algorithms	Training time (min)	Accuracy
RLPrompt	3100	87.7
TEMPERA	3208	90.5
POEM	**21**	**93.4**

Table 3. Comparison between POEM, RLPrompt and TEMPERA in terms of training time until convergence (minutes, the smaller the better) and accuracy of SST-2 dataset.

Algorithms	Accuracy
RLPrompt	52.6
TEMPERA	88
POEM	**93.4**

Table 4. Accuracy (60 iterations) of POEM, RLPrompt and TEMPERA on SST-2 dataset.

	SST-2	CR	COLA	AG News
Naive Action	91	91.4	68	79.6
Rank Action	**93.4**	**92.6**	**68.4**	**80.3**

Table 5. Ablation on action encoding. Average accuracy over 4 runs.

	SST-2	CR	COLA	AG News
Imbalanced labels	**93.5**	92.2	58.6	69.7
Balanced labels	93.4	**92.6**	**68.4**	**80.3**

Table 6. Average accuracy over 4 runs on imbalanced in-context examples selection.

Baselines. We note that complicated optimization baselines such as RLPrompt and TEMPERA have not been designed and applied to this task. Therefore, we compare POEM with several heuristic baselines. *0-shot* only includes the test query. *Random Ordering* randomly permute the examples. *Ascending Ordering* arranges in-context examples by increasing relevance, placing the most pertinent example closest to the test instance. In contrast, *Descending Ordering* positions examples from most to least relevant in relation to the test instance.

Experiment Setup. We use Llama-2-7b-chat, Llama-2-13b-chat, Llama-2-70b-chat [38] with default parameters. For all datasets, similar to classification tasks, we randomly sample 16 examples for training and another 16 examples to form an in-context dataset. For all baselines (except for 0-shot), we use $m = 4$ in-context examples for each test query. At test time, we select the number of nearest neighbors $k = 10$. Baselines that use ICL examples share the same example selection mechanism described in the Method section.

Results. We present our results in Tab. 2. Overall, it is clear that POEM shows superior results compared to Random and 0-shot baselines. This suggests that our approach is beneficial for optimizing ICL

the answers.

Datasets. We measure performance on two main tasks with four datasets, categorized as follows: **Commonsense Reasoning.** Hellaswag [43], PIQA [2]; **Question Answering.** TruthfulQA-mc1 [21], TriviaQA [12]. For more complex tasks such as Question Answering and Reading Comprehension, which involve multiple text fields, we employ a reranking approach based on the field containing the most relevant information. In cases like TruthfulQA, where certain fields like *type* and *category* lack semantic significance, we measure textual distances between examples using the *question* field, which typically encapsulates the query to be answered. For more detailed information on datasets and the fields used for retrieval, please refer to Appendices A.3 and A.4.

Parameter	Value	SST-2	CR	COLA	AG News
N	10	92.9	90.8	67.2	79.7
	60	**93.4**	**92.6**	**68.4**	**80.3**
	120	93.3	92.5	68.2	80.2
M	8	93	90.8	**68.7**	80
	12	93.1	91	**68.7**	80.2
	16	**93.4**	**92.6**	68.4	**80.3**
m	2	91.8	92.1	67.9	59.4
	4	**93.4**	**92.6**	**68.4**	**80.3**
	6	91.1	90.8	68.2	71.4
k	6	93.1	92.3	68.1	79.9
	10	**93.4**	**92.6**	**68.4**	**80.3**
	12	93.1	92.5	68.1	79.5

Table 7. Average accuracy across different N, M, m and k for few-shot text classification datasets.

Model	m	k	TruthfulQA	PIQA
	2	10	35.1	78.67
	4	10	**38.91**	**79.27**
Llama 2 7B	6	10	38.64	78.73
	2	10	36.33	78.94
	4	10	**40.95**	**79.43**
Llama 2 13B	6	10	40.27	79
	4	6	**39.46**	78.78
	4	10	38.91	**79.27**
Llama 2 7B	4	12	38.37	79
	4	6	40.54	**79.43**
	4	10	**40.95**	**79.43**
Llama 2 13B	4	12	40.41	78.84

Table 8. Accuracy across different m and k for general language understanding datasets.

example ordering. Compared to heuristic methods, in the context of 4 examples and for these datasets, POEM performs slightly better. This is because in these cases, the LLMs seem to favor the examples to be ordered descendingly. However, we note that POEM performs better in 9 out of 12 different settings compared to the Descending baseline. We also would like to highlight that POEM consistently shows up in the top 2 best baselines across all datasets. We acknowledge the challenges in replicating Llama's evaluation setup due to the unavailability of their system prompts. Consequently, to maintain simplicity and consistency, our study does not incorporate specific instructions into the prompts for LLMs. Instead, we apply uniform templates across all baselines. This approach ensures fair comparison between POEM and other baselines, thus showing the impact of in-context example reordering. For further details on the templates for this task, please see Tab. 10.

3.3 Ablation Studies

3.3.1 Analysis of Efficiency

We provide empirical evidence for our claim of POEM being fast and efficient. We compare performance and runtime on SST-2 dataset [35] between POEM, RLPrompt and TEMPERA. For a fair evaluation, all experiments were conducted with one identical GPU Nvidia A100. As shown in Tab. 3, the training time (until convergence) of POEM is approximately 150 times faster than that of TEMPERA and RLPrompt while achieving better accuracy. We also compare the performance of POEM with that of TEMPERA and RLPrompt after 60 iterations. The results in Table 4 reveal that POEM has attained a state of convergence, demonstrating the effectiveness of its learning algorithm within the given iteration frame. On the other hand, TEMPERA exhibits ongoing optimization efforts beyond the 60-iteration mark. This observation leads us to posit that TEMPERA's decent performance is attributable to its initial prompt construction methodology. RLPrompt, however, is still in the early stages of training, reflected by its near-random

accuracy. It is expected that RLPrompt will need more iterations to properly improve its prompt generation process.

3.3.2 Analysis of POEM's Components

Action encoding. We want to look further into how our action encoding contributes to our framework. To do this, we design a naive action encoding of a sequence of m examples. This results in an action space of $\frac{M!}{(M-m)!}$ instead of $m!$, where M and m are the number of total in-context examples and number of examples per prompt, respectively. As seen in Tab. 5, without our similarity-ranked encoding, performances on all classification tasks decrease. Notably, on SST-2 dataset, accuracy drops 2.4 %, from 93.4 to 91.

Imbalanced labels. For *few-shot classification* task, we aim to investigate how imbalanced in-context example labels can affect performance. Instead of employing our example selection strategy described in Section 2.2 to ensure an equal number of labels within a prompt, we now select the top-m closest examples to the test query regardless of their labels and use them to construct prompts. The results can be seen in Tab. 6. We can clearly see that for sentiment classification datasets (SST-2 and CR), POEM with imbalanced labels can still perform well. For datasets that require natural language understanding (NLU) like COLA, or those with more labels like AG News, the performances surprisingly dropped significantly. This perhaps indicates that there might be a strong bias coming from the imbalanced example selection for test queries.

3.3.3 Hyperparameter Sensitivity

Number of iterations. We report the results with different number of training iterations N. We vary the number of iterations to 10, 60 (default), and 120. The results in Table 7 suggest that more iterations may enhance **Memory Writing** by enabling POEM to find superior permutations. However, we set the default to 60 iterations to optimize the trade-off between training duration and convergence stability.

Size of in-context dataset $\mathcal{D}_{ic}$**.** We aim to further explore the impact of the in-context dataset size on performance. Intuitively, a broader selection of examples should yield more effective prompts. As indicated in Table 7, an increase in the in-context dataset size M results in a slight enhancement of performance.

Number of nearest neighbors k **in memory reading.** We provide results with different number of neighbors to study the robustness of POEM. With $k = 6$, $k = 10$ and $k = 12$, we obtain the results shown in Tab. 7 for *few-shot classification* datasets, and Tab. 8 for *general language understanding* datasets. It can be seen that POEM's performance holds for different k, which indices our method is relatively stable. We believe the reason behind this stability comes from the weighted sum formula in 10, meaning that the order is predominantly determined by test query's closest neighbors.

Number of examples per prompt. We ablate on the number of examples m used for demonstration in our algorithm. We choose the size of 2, 4, 6 for this analysis. For *few-shot classification*, as shown in Table 7, POEM performs optimally with 4 examples. A similar trend is observed for *general language understanding*, with results presented in Table 8. It is noteworthy that, in the original study [38], an increment in the number of examples from four to five also led to reduced accuracy. This implies that an overabundance of examples might introduce noise, thereby leading to incorrect model decisions.

Dataset	Natural Instructions
SST-2	"In this task, you are given sentences from movie reviews. The task is to classify a sentence as "great" if the sentiment of the sentence is positive or as "terrible" if the sentiment of the sentence is negative. "
IMDB	"In this task, you are given a review of movie. Your task is to classify given movie review into two categories: 1) great, and 2) terrible based on its content. "
CR	"In this task, you are given sentences from customer reviews. The task is to classify a sentence as "great" if the sentiment of the sentence is positive or as "terrible" if the sentiment of the sentence is negative. "
COLA	"You will be given a sentence. Check whether the sentence is grammatically correct and is meaningful. If the sentence is grammatically correct, then answer with "yes", otherwise answer with "no". "
AG News	"Classify the news articles into the categories of World, Sports, Business, and Technology. "
QNLI	"You are given two sentences. Your task is to determine if sentence 1 entails sentence 2. If sentence 1 entails sentence 2, answer with "yes", otherwise answer with "no". "
Boolq	"In this task you will be given a passage and a yes/no question based on the passage. You should answer the question using the information from the passage. "

Table 9. Natural instructions used for few-shot classification datasets.

Task	Templates
SST-2	"Does the following sentence have a {"great"} or {"terrible"} sentiment? {sentence}. Sentiment: {answer_choices[label]}"
IMDB	"Review: {sentence}. Sentiment: {answer_choices[label]}"
CR	"Review: {sentence}. Sentiment: {answer_choices[label]}"
COLA	"I'm copy-editing a story for publication. It has the following sentence in it: {sentence}. Does this sentence make sense and is it grammatically correct? Please answer yes or no. Answer: {answer_choices[label]}"
AG News	"What label best describes this news article? {sentence}. Answer: {answer_choices[label]}"
QNLI	"Passage: {passage}. After reading this passage, I have a question: {question}? yes or no? Answer: {answer_choices[label]}"
Boolq	"Passage: {passage}. After reading this passage, I have a question: {question}? True or False? {answer_choices[label]}"
Hellaswag	"{activity_label}: {ctx_a} {ctx_b}"
PIQA	"Question: {goal}. Answer:"
TruthfulQA	"Q: {question} A:"
TriviaQA	"Question: {question}? Answer:"

Table 10. Templates used for in-context examples for all datasets.

4 Related Works

Prompt Engineering. The traditional approach to using pre-trained LMs involves fine-tuning downstream datasets [8, 19], which involves extensive updates to model parameters. However, this method has shown limited success on downstream tasks. Another approach involves utilizing manual prompts to guide LLMs in performing NLP tasks without requiring additional training [4, 33, 34]. A different line of work in prompt engineering aims to develop instructional prompts, which offer task descriptions instead of fill-in-the-blank questions. In-context Learning [4, 10, 22, 26] achieves impressive performance by incorporating in-context demonstrations. However, these prompt engineering approaches are time-consuming and require manual tuning, which is not always feasible.

Prompt Optimization. Previous studies have also explored the application of RL for prompt optimization. Deng et al. (2022) propose using RL to directly generate prompts agnostic to specific queries [7]; however, the generated prompts may lack meaningfulness. Another previous RL editing method [44] allowed editing task descriptions and in-context examples. Yet, editing descriptions seldom aids optimization, and swapping examples can be time-consuming, potentially leading back to the initial state. Moreover, the method suits categorical tasks like classification. In the realm of continuous embedding space, gradients derived from LMs are employed to directly facilitate prompt optimization, a method also referred to as *soft prompt* tuning [20, 24]. However, due to their continuous nature, *soft prompts* pose challenges in comprehension [18] and lack reusability across diverse models because of disparities in latent embedding spaces [36]. With the expansion of LLMs in terms of both capacity and capabilities, there has emerged a new line of research that employs LLMs as prompt generators, prompt editors, or prompt scorers. [42] propose a method utilizing LLMs as prompt optimizers, where the generated prompts rely on the prior knowledge encoded within the LLMs. Zhou et al. (2022) utilize two distinct LLMs, one as a zero-shot instruction generator and the other as a scorer to optimize instructions [47]. However, these approaches share the drawback of relying heavily on the prior knowledge encoded within LLMs, necessitating high-quality LLMs, which can be resource-intensive. Furthermore, the variability in generated outputs by LLMs can be unpredictable and challenging to interpret.

Examplars retrieval and ordering in In-context learning. Research has demonstrated the significant influence of in-context example selection and arrangement on the performance of language models (LMs). For instance, Rubin et al. (2021) utilize a retriever to select in-context examples [32]. Additionally, Liu et al. (2021) propose a heuristic approach to ordering examples, ranking them based on the textual similarity between the test query and in-context examples [22]. Recently, Zhang et al. have applied RL to enable the swapping of in-context examples within their action space [44]. These studies underscore the impact of selecting and arranging in-context examples on downstream task performance. Compared to heuristics, RL solutions are theoretically guaranteed. Unfortunately, existing RL-based methods are complicated and slow during training. Our study is the first RL-based prompt optimization method that is both simple and efficient, demonstrating superior performance compared to existing counterparts.

5 Discussion

We have introduced POEM, a novel approach to prompt optimization within the RL paradigm. By reordering few-shot examples using episodic memory, POEM significantly boosts the performance of LLMs on a range of NLP tasks. Our method outperforms existing techniques in few-shot classification on the various datasets and shows advancements over heuristic baselines in general language understanding tasks. The synergy between NLP and RL in POEM highlights the potential for future innovations in test-time prompt optimization algorithms, which could be pivotal for practical applications.

References

[1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[2] Y. Bisk, R. Zellers, J. Gao, Y. Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439, 2020.

[3] C. Blundell, B. Uria, A. Pritzel, Y. Li, A. Ruderman, J. Z. Leibo, J. Rae, D. Wierstra, and D. Hassabis. Model-free episodic control. *arXiv preprint arXiv:1606.04460*, 2016.

[4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[5] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.

[6] C. Clark, K. Lee, M.-W. Chang, T. Kwiatkowski, M. Collins, and K. Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, 2019.

[7] M. Deng, J. Wang, C.-P. Hsieh, Y. Wang, H. Guo, T. Shu, M. Song, E. Xing, and Z. Hu. Rlprompt: Optimizing discrete text prompts with reinforcement learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3369–3391, 2022.

[8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. URL https://api.semanticscholar.org/CorpusID:52967399.

[9] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.

[10] T. Gao, A. Fisch, and D. Chen. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, 2021.

[11] M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, page 168–177, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138881. doi: 10.1145/1014052.1014073. URL https://doi.org/10.1145/1014052.1014073.

[12] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, 2017.

[13] D. Khashabi, X. Lyu, S. Min, L. Qin, K. Richardson, S. Welleck, H. Hajishirzi, T. Khot, A. Sabharwal, S. Singh, et al. Prompt waywardness: The curious case of discretized interpretation of continuous prompts. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3631–3643, 2022.

[14] H. Le and S. Venkatesh. Neurocoder: General-purpose computation using stored neural programs. In *International Conference on Machine Learning*, pages 12204–12221. PMLR, 2022.

[15] H. Le, T. Karimpanal George, M. Abdolshah, T. Tran, and S. Venkatesh. Model-based episodic memory induces dynamic hybrid controls. *Advances in Neural Information Processing Systems*, 34:30313–30325, 2021.

[16] H. Le, M. Abdolshah, T. K. George, D. Kim, D. Nguyen, and S. Venkatesh. Episodic policy gradient training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 7317–7325, 2022.

[17] M. Lengyel and P. Dayan. Hippocampal contributions to control: the third way. *Advances in neural information processing systems*, 20, 2007.

[18] B. Lester, R. Al-Rfou, and N. Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, 2021.

[19] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for language generation, translation, and comprehen-sion. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.

[20] X. L. Li and P. Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, 2021.

[21] S. Lin, J. Hilton, and O. Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, 2022.

[22] J. Liu, D. Shen, Y. Zhang, W. B. Dolan, L. Carin, and W. Chen. What makes good in-context examples for gpt-3? In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114, 2022.

[23] X. Liu, K. Ji, Y. Fu, W. Tam, Z. Du, Z. Yang, and J. Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, 2022.

[24] X. Liu, Y. Zheng, Z. Du, M. Ding, Y. Qian, Z. Yang, and J. Tang. Gpt understands, too. *AI Open*, 2023.

[25] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

[26] Y. Lu, M. Bartolo, A. Moore, S. Riedel, and P. Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, 2022.

[27] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts. Learning word vectors for sentiment analysis. In D. Lin, Y. Matsumoto, and R. Mihalcea, editors, *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL https://aclanthology.org/P11-1015.

[28] P. Pezeshkpour and E. Hruschka. Large language models sensitivity to the order of options in multiple-choice questions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017, 2024.

[29] A. Prasad, P. Hase, X. Zhou, and M. Bansal. Grips: Gradient-free, edit-based instruction search for prompting large language models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3845–3864, 2023.

[30] R. Pryzant, D. Iter, J. Li, Y. Lee, C. Zhu, and M. Zeng. Automatic prompt optimization with "gradient descent" and beam search. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, 2023.

[31] N. Reimers and I. Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 2019.

[32] O. Rubin, J. Herzig, and J. Berant. Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, 2022.

[33] V. Sanh, A. Webson, C. Raffel, S. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, A. Raja, M. Dey, et al. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*, 2021.

[34] T. Schick and H. Schütze. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, 2021.

[35] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.

[36] Y. Su, X. Wang, Y. Qin, C.-M. Chan, Y. Lin, H. Wang, K. Wen, Z. Liu, P. Li, J. Li, et al. On transferability of prompt tuning for natural language processing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3949–3969, 2022.

[37] T. Sun, Y. Shao, H. Qian, X. Huang, and X. Qiu. Black-box tuning for language-model-as-a-service. In *International Conference on Machine Learning*, pages 20841–20855. PMLR, 2022.

[38] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*,

2023.

[39] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, 2018.

[40] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. *Advances in neural information processing systems*, 32, 2019.

[41] Z. Wu, S. Wang, J. Gu, R. Hou, Y. Dong, V. V. Vydiswaran, and H. Ma. Idpg: An instance-dependent prompt generation method. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5507–5521, 2022.

[42] C. Yang, X. Wang, Y. Lu, H. Liu, Q. V. Le, D. Zhou, and X. Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=Bb4VGOWELI.

[43] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

[44] T. Zhang, X. Wang, D. Zhou, D. Schuurmans, and J. E. Gonzalez. Tempera: Test-time prompt editing via reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2022.

[45] X. Zhang, J. Zhao, and Y. LeCun. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.

[46] Z. Zhao, E. Wallace, S. Feng, D. Klein, and S. Singh. Calibrate before use: Improving few-shot performance of language models. *ICML*, 2021.

[47] Y. Zhou, A. I. Muresanu, Z. Han, K. Paster, S. Pitis, H. Chan, and J. Ba. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*, 2023.

A Appendix

A.1 Additional results

We provide additional results for comparing POEM with heuristic baselines in Tab. 12. We note that our method consistently surpasses heuristic baselines. We also notice that the heuristic trend varies between different datasets.

A.2 Training details

We provide the training details in Tab. 11. We also include the parameters for LLMs such as top p and temperature.

Training details	
Minibatch size	16
Number of iterations (N)	60
Number of examples per label (g)	16
Number of in-context examples in total (M)	16
Number of training samples for non-categorical tasks	80
Number of in-context examples per prompt (m)	4
Number of nearest neighbors (k)	10
Positive lambda coefficient (λ_1)	2.0
Negative lambda coefficient (λ_2)	1.8
LLMs' top p	0.9
LLMs' temperature	0.6

Table 11. Hyperparameters used in our experiments.

A.3 Dataset details

We provide information for all datasets in Tab. 13 and Tab. 14. For Finetuning, we use standard finetuning of the RoBERTa model from huggingface for 100 epochs, a learning rate of 0.0003 and the optimizer of Adam.

A.4 Prompting details

For datasets that have more than one text field, we provide the fields used for retrieve nearest neighbor in 15.

	SST-2	Boolq	CR	AG News	IMDB	QNLI	COLA	**Average**
Ascending Ordering	92.5	64.8	<u>92.4</u>	**80.9**	<u>90.8</u>	54.0	66.7	<u>77.4</u>
Descending Ordering	<u>92.7</u>	<u>65.9</u>	86.1	74.5	54.6	**55.3**	<u>67.9</u>	71
POEM (Ours)	**93.4(0.2)**	**66.1(0.1)**	**92.6(0.2)**	<u>80.3(1.1)</u>	**90.9(0.1)**	<u>54.3(0.2)</u>	**68.4(0.4)**	**78**

Table 12. Accuracy and standard derivation of POEM over 4 seeds versus advanced heuristic baselines on few-shot classification. The highest accuracy is **bolded** and the second-highest accuracy is <u>underlined</u>. The last column shows the average accuracy across all datasets in this table.

Dataset details							
Category	Dataset	Number of labels	$|\mathcal{D}_{train}|$	$|\mathcal{D}_{test}|$	Type	Manual Prompt	Verbalizers
Single sentence	SST-2	2	32	1.8k	Sentiment	$\langle S \rangle$ It was $[mask]$.	great, terrible
	IMDB	2	32	2k	Sentiment	It was $[mask]$.	great, terrible
	CR	2	32	2k	Sentiment	It was $[mask]$.	great, terrible
	COLA	2	32	1.1k	NLU	It was $[mask]$.	yes, no
	AG News	4	64	7.6k	Topic	$[mask]$ News: $\langle S \rangle$	World, Sports, Business, Tech
Sentence-pair	QNLI	2	32	9.8k	NLI	$\langle S_1 \rangle$? $[mask]$,$\langle S_2 \rangle$	yes, no
	Boolq	2	32	3.3k	MRC	$\langle S_1 \rangle$? $[mask]$,$\langle S_2 \rangle$	yes, no

Table 13. Classification dataset details

Dataset details				
Dataset	$|\mathcal{D}_{train}|$	$|\mathcal{D}_{ic}|$	$|\mathcal{D}_{test}|$	Type
Hellaswag	16	16	10k	Commonsense Reasoning
PIQA	16	16	1.8k	Commonsense Reasoning
TruthfulQA	16	16	0.7k	Question Answering
TriviaQA	16	16	18k	Question Answering

Table 14. General Language Understanding dataset details

Task	Dataset	Field name
Classification	Boolq	{question} + {passage}
	QNLI	{text1} + {text2}
General Language Understanding	Hellaswag	{question}
	PIQA	{goal}
	TruthfulQA	{question}
	TriviaQA	{question}

Table 15. Selected fields for nearest neighbor retrieval for datasets with multiple fields