

Pre-trained Language Models Improve the Few-shot Prompt Ability of Decision Transformer

Yu Yang
yu.yang@duke.edu
Duke University

Pan Xu
pan.xu@duke.edu
Duke University

Abstract

Decision Transformer (DT) has emerged as a promising class of algorithms in offline reinforcement learning (RL) tasks, leveraging pre-collected datasets and Transformer’s capability to model long sequences. Recent works have demonstrated that using parts of trajectories from training tasks as prompts in DT enhances its performance on unseen tasks, giving rise to Prompt-DT methods. However, collecting data from specific environments can be both costly and unsafe in many scenarios, leading to suboptimal performance and limited few-shot prompt abilities due to the data-hungry nature of Transformer-based models. Additionally, the limited datasets used in pre-training make it challenging for Prompt-DT type of methods to distinguish between various RL tasks through prompts alone. To address these challenges, we introduce the Language model-initialized Prompt Decision Transformer (LPDT), which leverages pre-trained language models for meta-RL tasks and fine-tunes the model using Low-rank Adaptation (LoRA). We further incorporate prompt regularization to effectively differentiate between tasks based on prompt feature representations. Our approach integrates pre-trained language model and RL tasks seamlessly. Extensive empirical studies demonstrate that initializing with a pre-trained language model significantly enhances the performance of Prompt-DT on unseen tasks compared to baseline methods.

1 Introduction

In many sequential decision-making applications such as robotic manipulation and autonomous driving (Sinha et al., 2022; Kumar et al., 2021), it can be expensive or even unsafe for agents to learn through trial-and-error with the environment. Offline reinforcement learning (RL) methods (Levine et al., 2020) have emerged as a powerful paradigm for optimizing agent policies without directly interacting with the environment. They leverage pre-collected datasets obtained from a set of behavior policies instead of online interactions to learn an optimal policy. Among these Offline RL methods, Decision Transformer (DT) (Chen et al., 2021) has become popular for offline RL tasks due to its scalability with computation and data and stability in training. DT models a goal-conditioned policy using a Transformer network, solving a sequence-prediction problem in a supervised learning manner. More specifically, DT formulates decision-making as sequence generation over pre-collected trajectories using the powerful Transformer architecture (Vaswani et al., 2017). DT models the states, actions and return-to-go from the RL trajectory as the word tokens in Natural Language Processing (NLP) tasks and then generates the actions conditioned on the return goals.

Compared with traditional dynamic programming-based offline RL methods (Kumar et al., 2019; Fujimoto et al., 2019; Kumar et al., 2020) that heavily rely on the Markov Decision Process (MDP) assumption of the environment, Decision Transformer can utilize entire trajectory histories to predict the next action, making them more applicable in partially observable environments where all past information must be incorporated in decision-making (Kaelbling et al., 1998; Ni et al., 2024). Furthermore, the supervised learning nature of DTs enhances the stability and scalability in the

training process compared to dynamic programming algorithms based on Bellman equations (Chen et al., 2021; Janner et al., 2021; Zheng et al., 2022). Another advantage of Transformers is their few-shot generalization ability (Brown et al., 2020; Achiam et al., 2023). Based on their remarkable few-shot generalization capability, a prompt-based framework has been proposed and proven effective for adapting to new tasks in NLP (Brown et al., 2020; Li & Liang, 2021). In this paradigm, the prompt, containing useful information about the task, is inputted as a prefix to the model for identifying the environments. Previous works have demonstrated that DTs also exhibit good generalization ability for unseen tasks. Prompt-DT (Xu et al., 2022) leverages parts of trajectories from datasets as prompts to encapsulate task information. The method is trained on these trajectories and corresponding prompts, then tested on unseen tasks with few-shot demonstrations as prompts.

However, existing Prompt-DT methods (Xu et al., 2022; Hu et al., 2023; 2024) require significant amounts of data for pre-training due to the data-hungry nature of Transformers (Brown et al., 2020; Achiam et al., 2023). Offline RL datasets are often small and insufficient to fully unleash the few-shot prompt learning capability of Transformers. Collecting large amounts of RL trajectories for pre-training powerful Decision Transformers is challenging. Inspired by the broad success of large language models in NLP, recent works (Li et al., 2022; Reid et al., 2022; Shi et al., 2023) have shown the potential of such models to provide effective initial weights for decision-making tasks. However, these works do not directly demonstrate few-shot abilities due to a lack of multi-task training and prompt guidance. Language initialization in these works provides pre-knowledge and helps alleviate the need of huge datasets. Therefore, we aim to explore the use of pre-trained language models to initialize Prompt-DT methods and reduce the dependency on large datasets for training.

In this work, we propose a novel framework, Language model-initialized Prompt Decision Transformer (LPDT), that utilizes pre-trained language model initialization to improve the few-shot prompt ability of Decision Transformer. Our approach initializes the model with pre-trained language models, incorporating pre-existing knowledge that might benefit downstream RL tasks. We combine this pre-trained knowledge with domain-specific knowledge from multi-task RL by using LoRA (Hu et al., 2021), a parameter-efficient fine-tuning method, to fine-tune the pre-trained model on a multi-task RL offline dataset using prompts. Furthermore, distinguishing different testing environments is vital for multi-task or meta RL. Therefore, we introduce prompt regularization methods to help the pre-trained Decision Transformer distinguish different RL tasks, guiding action generation under new, unseen task-specific prompt representations. A more detailed illustration of the model structure and our training paradigm is provided in Figure 1. We conduct extensive experiments to assess the capability of our proposed framework in MuJoCo control environments (Fu et al., 2020) and Meta World ML1 tasks (Yu et al., 2020). Our method outperforms baselines in terms of cumulative rewards on unseen tasks. Our contributions are summarized as follows.

- We propose a framework named LPDT to improve the few-shot prompt ability of Decision Transformer. This framework involves leveraging the language model as the initialization of DT and imposing both supervised and unsupervised prompt regularization. LPDT demonstrates improved few-shot prompt capabilities with pre-trained language knowledge in multi-task RL.
- We utilize Low-Rank Adaptation (LoRA) and an additional prompt regularization method to combine pre-trained knowledge with domain-specific RL task knowledge. LoRA allows efficient fine-tuning by adapting only a small subset of parameters, while both the supervised and unsupervised prompt regularization enhances the model’s ability to distinguish task information contained in the prompts.
- Through extensive experiments on MuJoCo control and Meta World ML1, we demonstrate the advantages of LPDT compared to baselines. Our results show that LPDT significantly outperforms existing models in performance under full and limited datasets, highlighting its potential for real-world applications.

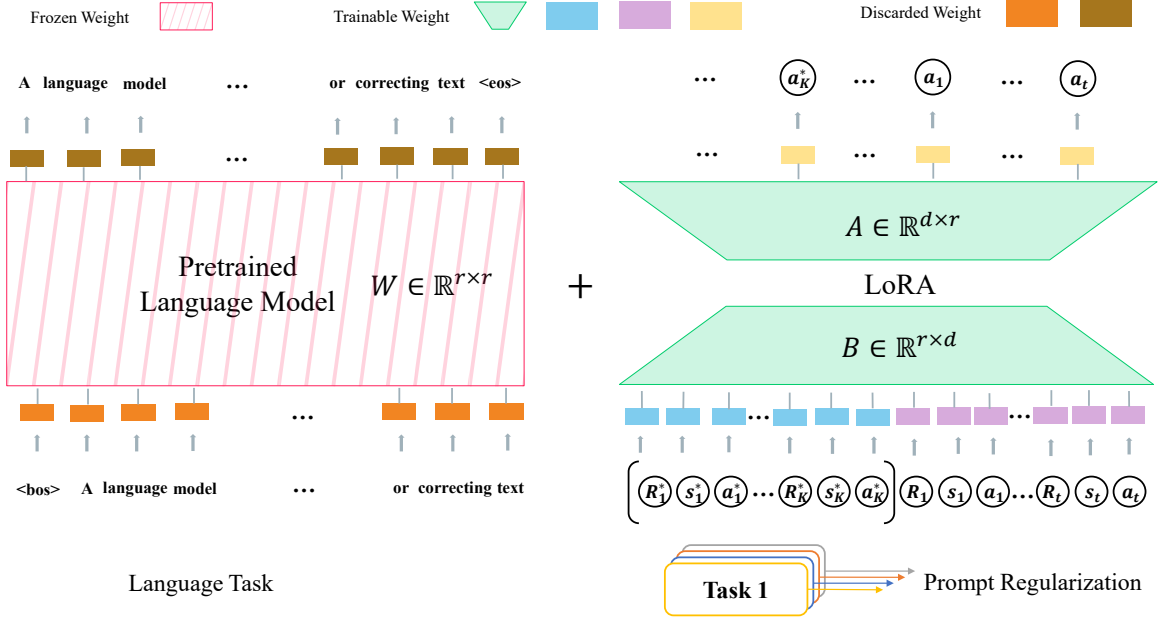


Figure 1: Overview of LPDT. We first initialize our algorithm using a pre-trained language model such as DistilGPT2 (Sanh et al., 2019). The pre-trained language model is trained on a large corpus of data using the causal language modeling objective which is to predict the next-token paradigm. Our method LPDT replaces the word embedding layers with linear layers to fully learn and capture the features of RL trajectory tokens. We fine-tune our model using parameter-efficient methods like Low-Rank Adaptation (LoRA). Specifically, we freeze the initial weights of the language models and update only the LoRA weights. The input to our approach consists of prompts accompanied by training trajectories from the same tasks. Unlike traditional models that predict word tokens, our method predicts action tokens. Additionally, we incorporate prompt regularization over the input prompts. This is achieved by introducing an additional loss on the prompt embeddings, which helps LPDT distinguish between different environments. More technical details of our method are presented in Section 3.

2 Preliminary

2.1 Offline Reinforcement Learning

Reinforcement learning problems are usually formulated as a Markov Decision Process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, d_0, \mathcal{R}, \gamma)$, where \mathcal{S} represents the set of states $s \in \mathcal{S}$, \mathcal{A} represents the set of actions $a \in \mathcal{A}$, \mathcal{T} is the transition distribution in the form $\mathcal{T}(s_{t+1}|s_t, a_t)$, d_0 is the distribution of initial states s_0 , $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $r_t = \mathcal{R}(s_t, a_t)$ is the reward at timestep t , and $\gamma \in (0, 1)$ is the discount factor. The objective is to find a policy π that maximizes the expected cumulative rewards $J(\pi)$:

$$J(\pi) = \mathbb{E}_{s_0 \sim d_0(\cdot), a_t \sim \pi(\cdot|s_t), s_{t+1} \sim \mathcal{T}(\cdot|s_t, a_t)} \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t) \right]. \quad (1)$$

There are many RL algorithms proposed to solve these MDP problems via online interactions with the environment. However, online RL algorithms are not feasible in some scenarios such as robotic manipulation and autonomous driving where the interactions with the environment are associated with high computational cost and risks. Offline RL methods (Levine et al., 2020) become popular in these scenarios. Different from the online setting, the agent has access to a dataset \mathcal{D} containing trajectories collected by a behavior policy instead of access to the environment. The agent is expected to find the optimal policy using only the offline dataset \mathcal{D} , without interacting with the

environment itself. Among various offline RL methods (Levine et al., 2020; Kumar et al., 2019; Fujimoto et al., 2019; Kumar et al., 2020), Decision Transformer (Chen et al., 2021) which leverages the Transformer (Vaswani et al., 2017) to predict the next action conditioned on the past trajectory is drawing increasing attention. Transformer is first proposed by Vaswani et al. (2017), which has been extensively used in natural language processing (NLP) and computer vision (CV). Recently, Transformer has also been increasingly studied in reinforcement learning using the sequence modeling paradigm (Chen et al., 2021; Furuta et al., 2021; Xu et al., 2022; Janner et al., 2021). In DT, the trajectories $\{s_0, a_0, r_0, s_1, a_1, r_1, \dots, s_T, a_T, r_T\}$ in the offline dataset \mathcal{D} are reformulated and modeled as a sequence generation problem via self-supervised learning paradigm.

2.2 Prompt Decision Transformer

Decision Transformer leverages the powerful Transformer model to predict the action via the sequence generation paradigm. Recent advances in NLP (Achiam et al., 2023; Puri & Catanzaro, 2019) demonstrate that Transformer enjoys impressive few-shot learning capabilities when it is pre-trained on vast language datasets and then generates the word tokens under a specific prompt in the testing phase which is used to describe the language tasks. Similarly, Decision Transformer can also be extended to use prompt information to improve its generalization ability in unseen tasks during testing. Xu et al. (2022) introduced Prompt-DT, which leverages the Decision Transformer architecture to model the RL trajectories in multi-task environments and make decisions in unseen tasks based on the prompt framework, achieving few-shot learning ability in offline RL setting. Different from the prompt-based methods in NLP tasks, Prompt-DT utilizes the sampled short trajectories from the specific task in the offline dataset as the prompt. These prompts also contain environment information, which can help the model distinguish the tasks.

In the offline dataset \mathcal{D} , we have the trajectories τ . The prompts and training trajectories are both from this dataset. Different from the offline RL setting, the rewards in the training trajectories of Decision Transformer are replaced by the return-to-go which are denoted as the $R_i = \sum_{t=i}^T r_t$. The prompt is the short trajectory which can also be denoted as the tuples of state s^* , action a^* , and return-to-go R^* . It captures essential and concise trajectory information to aid task identification without prior knowledge of the RL tasks. During the training stage, we utilize the offline RL dataset \mathcal{D} which contains multiple RL tasks denoted as $T_i \in T^{train}$. The input of Prompt-DT is a concatenation of the prompt and the training trajectory, denoted by τ_i^* and τ_i respectively. Specifically, we denote the prompt τ_i^* as the following sequence

$$\tau_i^* = (R_{i,1}^*, s_{i,1}^*, a_{i,1}^*, \dots, R_{i,K^*}^*, s_{i,K^*}^*, a_{i,K^*}^*), \quad (2)$$

where K^* is the length of the prompt. These trajectory prompts are shorter than the horizon of the tasks, thereby providing crucial guidance without enabling complete task imitation. Beyond the prompt, the training trajectories of DT can be denoted as

$$\tau_i = (R_{i,1}, s_{i,1}, a_{i,1}, \dots, R_{i,K}, s_{i,K}, a_{i,K}), \quad (3)$$

where K is the length of the training trajectories. Consequently, we have the input vector τ_i^{input} defined as

$$\tau_i^{input} = [\tau_i^*, \tau_i] = (R_{i,1}^*, s_{i,1}^*, a_{i,1}^*, \dots, R_{i,K^*}^*, s_{i,K^*}^*, a_{i,K^*}^*, R_{i,1}, s_{i,1}, a_{i,1}, \dots, R_{i,K}, s_{i,K}, a_{i,K}). \quad (4)$$

Besides, we denote the partial trajectory from the timestep 1 to timestep t as $\tau_{i,1<t}^{input}$. Then the learning objective of Prompt-DT can be formulated as the following maximum likelihood estimation:

$$L_{PDT} = \mathbb{E}_{\tau_i^{input} \sim T_i} [\sum_{t=1}^K -\log M_\theta(\hat{a}_{i,t} | \tau_i^*, \tau_{i,1<t}^{input}, R_{i,t}, s_{i,t})]. \quad (5)$$

where M_θ denotes the Prompt-DT model with the parameter θ . In practical implementations, we often use the mean squared error loss instead, which aims to predict the future action $\hat{a}_{i,t}$ given the history trajectory and current state by minimizing the following loss function.

$$L_{PDT} = \mathbb{E}_{\tau_i^{input} \sim T_i} [1/K \sum_{t=1}^K (a_{i,t} - \hat{a}_{i,t})^2]. \quad (6)$$

The training procedure of Prompt-DT is to autoregressively generate the action conditioned on the current state, return-to-go, past trajectory and sampled prompt.

3 The Proposed Framework

We propose Language model-initialized Prompt Decision Transformer (LPDT), a novel and effective framework to incorporate powerful pre-trained language models into Decision Transformers to improve their few-shot learning abilities. We also leverage an additional prompt regularization over the prompts during fine-tuning to better identify tasks. Figure 1 illustrates the overview of our method. At a high level, LPDT incorporates several key components:

- **Language model initialization for Prompt-DT:** We first use a pre-trained language model, such as DistilGPT2 (Sanh et al., 2019; Radford et al., 2019), as the initialization for our Decision Transformer. This design ensures compatibility with the Prompt-DT paradigm, where prompts from various tasks are appended to the input sequence.
- **Parameter-efficient fine-tuning on RL tasks:** We adopt Low-Rank Adaptation (LoRA) (Hu et al., 2021) to fine-tune a low-rank residual matrix while keeping the original weight matrix of the language model fixed throughout the learning process. This approach significantly reduces the number of parameters compared to standard full fine-tuning of the large language models.
- **Prompt regularization with supervised and unsupervised objectives:** We incorporate additional regularization over the prompt embeddings to better identify tasks. Specifically, we employ loss functions derived from both supervised and unsupervised learning techniques to fully utilize task-related information from prompts, thereby preventing the language model from overfitting specific tasks.

We discuss these techniques in detail in the rest of this section. At the end of this section, built on these components, we present our learning algorithm in Algorithm 1.

3.1 Language model initialization for Prompt-DT

The first step in our LPDT framework is to use pre-trained language models as the initialization. Recent advances in large language models have demonstrated that these models possess strong few-shot learning abilities. With task-specific information such as prompts for translation or answering questions, language models can generate relevant outputs. We adapt these language models to RL tasks, such as MuJoCo controls, to reduce the demand for large datasets by leveraging pre-trained knowledge that may have relevance to downstream RL tasks. In this work, we use the DistilGPT2 (Sanh et al., 2019; Radford et al., 2019) as the initial weights, which is a pre-trained model with 82 million parameters and is faster and lighter than the original GPT-2 (Radford et al., 2019). The common next-token prediction objective of GPTs can be formulated as

$$L_{LM} = \sum_{i=1}^{t-1} -\log(M_{\theta^*}(w_{i+1} | w_1, \dots, w_i)), \quad (7)$$

where M_{θ^*} is the language model and w_i represents the word token. To make the language model compatible with the RL sequence prediction tasks in DT, we follow previous work (Shi et al., 2023) to replace the word token embedding input and output layers with linear layers, which are trainable for specific RL tasks. Moreover, in the input layer, we also need to incorporate a prefix as the prompt from sequences chosen from trajectories rolled out in the offline dataset. The embedding layers are represented by the different color blocks in Figure 1 for the prompt and input sequence, respectively. Apart from the input and output embedding layers, the entire intermediate layers of the language model will be frozen during the training on RL tasks.

3.2 Parameter-efficient fine-tuning on RL tasks

To adapt the language models to specific RL tasks, we add a low-rank adaptation of the frozen weights of the language model and update it using parameter-efficient methods like LoRA (Hu et al.,

2021). Specifically, LoRA utilizes two low-rank matrices to represent the weight matrix, significantly reducing the number of parameters. This process can be formulated as $W = W_0 + \Delta W = W_0 + AB$, where $W \in \mathbb{R}^{d \times k}$ is the weight of our model, $W_0 \in \mathbb{R}^{d \times k}$ is the frozen weight inherited from the language model, and $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times k}$ are low-rank matrices. In this way, we avoid fully fine-tuning the language model and make our method scalable to large language models, where only a small number of parameters from the low-rank matrix ΔW need to be updated.

3.3 Prompt regularization with supervised and unsupervised objectives

Previous works such as Prompt-Tuning DT (Hu et al., 2023) and Prompt Diffuser (Hu et al., 2024) aim to tune the prompt during testing on unseen tasks. These methods seek to optimize the prompt during the testing phase separately, allowing the model to distinguish the current task during testing. However, they are not always effective and can result in inferior performance when tasks are too similar. To address this challenge and achieve improved performance on testing tasks, we incorporate a task identification procedure into the training process. To effectively distinguish tasks by their prompts, we adopt additional regularization on the training loss over the prompt embeddings, termed prompt regularization.

Since our model is built upon Prompt-DT, we use the loss function for Prompt-DT defined in (6) as the base loss function, and then incorporate a prompt regularization term. The final loss function of our method is as follows.

$$L_{\text{total}} = \mathbb{E}_{\tau_i^{\text{input}} \sim T_i} [1/K \sum_{t=1}^K (a_{i,t} - \hat{a}_{i,t})^2] + \lambda L_\phi, \quad (8)$$

where L_ϕ is the loss for the prompt regularization which we will specify in the rest of this section, and λ is the hyperparameter for prompt regularization.

In particular, we propose two practical implementations of prompt regularization based on supervised learning and unsupervised learning methods respectively.

Supervised learning-based prompt regularization. In this approach, we add a classifier head to the output of the prompt encoder. We use the task ID from the dataset as the label to help the prompt encoder learn a more meaningful embedding that can easily distinguish different task environments. We adopt cross-entropy as the loss function. We formulate L_ϕ as:

$$L_\phi^{\text{classifier}} = - \sum_i y_i \log(\hat{y}_i), \quad (9)$$

where y_i is the true task label which means the task ID and \hat{y}_i is the predicted probability for the task which comes from the prompt τ_i^* .

Unsupervised learning-based prompt regularization. When task IDs are unknown, the supervised method may not be feasible. Therefore, we also propose an unsupervised learning method to learn the prompt representation. From an information theory perspective, the ideal prompt encoder should aim to maximize the mutual information between the prompt representation and the tasks. We use the InfoNCE objective (Oord et al., 2018) to calculate the loss over the prompt. We formulate L_ϕ as:

$$L_\phi^{\text{InfoNCE}} = -\mathbb{E} \left[\log \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}{\sum_{k=1}^N \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k)/\tau)} \right], \quad (10)$$

where \mathbf{z}_i and \mathbf{z}_j are the encoded representations of the prompts τ_i^* and τ_j^* respectively. The term $\text{sim}(\mathbf{z}_i, \mathbf{z}_j)$ denotes the similarity function (e.g., cosine similarity) between \mathbf{z}_i and \mathbf{z}_j .

4 Experiments

In this section, we conduct experiments to evaluate the few-shot generalization ability of our proposed method LPDT. We evaluate the performance of LPDT on MuJoCo control tasks (Fu et al., 2020)

Algorithm 1: Language model-initialized Prompt Decision Transformer (LPDT)

Input: Pre-trained language model weights θ^* , training datasets \mathcal{D} with prompts and trajectories, prompt regularization hyperparameter λ

Output: Language Initialized Prompt Decision Transformer

- 1 Initialize Decision Transformer with pre-trained language model weights (e.g., DistillGPT2);
- 2 Replace the input and output embedding layers of the language model with linear layers;
- 3 **for** *each epoch* **do**
- 4 **for** *each batch in training dataset* **do**
- 5 Extract prompts τ_i^* and trajectories τ_i from the batch;
- 6 Encode prompts τ_i^* using the prompt encoder $\phi(\tau_i^*)$;
- 7 Transform trajectories τ_i using DT with prompt embeddings $\phi(\tau_i^*)$;
- 8 Compute the Prompt-DT loss L_{PDT} over the input trajectories τ_i^{input} by (6); Compute the prompt regularization loss L_ϕ using the supervised learning (9) or unsupervised learning loss (10);
- 9 Combine the Prompt-DT loss and prompt regularization loss: $L_{\text{total}} = L_{\text{PDT}} + \lambda L_\phi$;
- 10 Backpropagate the combined loss and update parameters using parameter-efficient methods like LoRA;
- 11 Freeze the remaining weights and update only low-rank matrices A and B :

$$W = W_0 + \Delta W = W_0 + AB$$
- 12 **return** Language model-initialized Prompt Decision Transformer;

and Meta World (Yu et al., 2020) with the episode accumulated reward as the evaluation metric. We also evaluate the prompt ability of LPDT with the smaller dataset sizes. Our experiments aim to address the following questions: (1) Can LPDT with language model initialization achieve better performance compared with Prompt-DT and other baselines? (2) Does the prompt regularization help the model distinguish the tasks and improve the prompt learning ability of LPDT? (3) Does the language-initialized Transformer model contain the knowledge of the unseen RL tasks and help improve the performance on a smaller size of data?

4.1 Implementation

In the empirical study, we implement our LPDT method with DistilGPT2 as the language initialization. The initialization language model weight comes from the Huggingface. The DistilGPT2 contains 82 million parameters which is lighter and more efficient than GPT2 with 124 million parameters. DistilGPT2 is pre-trained on the openwebtext (Puri & Catanzaro, 2019) and is distilled by Sanh et al. (2019). During the fine-tuning stage, we follow the same hyperparameters for Prompt-DT (see Appendix C for detail). We also leverage the LoRA to highly reduce the parameters trained. For the prompt regularization, we use MLP to further encode the prompt embedding. For the supervised version of prompt regularization defined in (9), we directly use the logits from the MLP to compute the cross-entropy loss and refer to the method as LPDT-Classifier. For the unsupervised version of prompt regularization defined in (10), we calculate the similarity matrix through the cosine similarity based on the logits from the MLP and refer to it as LPDT-InfoNCE.

4.2 Datasets and Tasks

In this work, we evaluate the performance of our proposed approach on MuJoCo controls and Meta World, which are commonly used in existing Prompt-DT type of methods (Xu et al., 2022; Hu et al., 2023; 2024), namely, Cheetah-dir, Cheetah-vel, Ant-dir, Meta-World reach-v2 and MW pick-place-v2. In Cheetah-dir, there are two tasks with goal directions as forward and backward, where the reward function promotes high velocity along the goal direction. The training and testing phases

both include the two tasks. Similar to Cheetah-dir, Ant-dir also segments the tasks by directions. There are 50 tasks in Ant-dir with different goal directions uniformly sampled in 2D space. The tasks are split into 45 training tasks and 5 testing tasks. The ant is also rewarded with high velocity along the goal direction. Different from segmenting the tasks by direction, Cheetah-vel penalizes the agent through the l_2 errors with the target velocities sampled from the velocity interval. There are 40 tasks with different goal velocities where 35 tasks are training tasks and 5 tasks are testing tasks. Except for the MuJoCo control meta-RL tasks, we also test our approach on Meta World (Yu et al., 2020) which is an open benchmark for meta-RL and multi-task learning. In this work, we evaluate our approach on Meta-World reach-v2 and Meta-World pick-place-v2. The objective of reach-v2 is to control the robot to reach the target position in 3D positions and pick-place-v2 is to grasp the object. Each task has a different goal position.

We utilize the dataset and settings from the Prompt-DT paper (Xu et al., 2022). To be specific, the datasets of Cheetah-dir and Ant-dir come from the replay buffer of Soft Actor-Critic (Haarnoja et al., 2018) and the dataset of Cheetah-vel comes from TD3 (Fujimoto et al., 2018). For Meta-World reach-v2 and Meta-World pick-place-v2, we collected the dataset through the expert policies provided in the open benchmark.

4.3 Baselines

We compare the few-shot generalization ability of our proposed LPDT with baseline algorithms. For each method, we compare the performance based on the accumulated reward. The baselines we choose include Prompt-DT (Xu et al., 2022), Prompt-Tuning DT (Hu et al., 2023), and Prompt Diffuser (Hu et al., 2024). Prompt-DT is the first method to utilize the prompt to guide Decision Transformer in testing with the few-shot demonstrations. Prompt-DT directly uses the prompt without any additional fine-tuning process when testing. Prompt-Tuning DT is based on Prompt-DT and utilizes prompt tuning methods when testing on the unseen task. Several prompt tuning techniques are used to tune the prompt to the specific target environment using preference ranking. Prompt Diffuser extends the prompt tuning method by leveraging diffusion models to generate high-quality prompts to improve the few-shot demonstration guidance. Beyond these baselines, Multi-Task Decision Transformer (MT-ORL) (Chen et al., 2021) is mentioned in Prompt-DT and Soft-Prompt (Lester et al., 2021) is described in Prompt Diffuser. So we do not demonstrate them in our experiments. For HDT (Xu et al., 2023b), it utilizes the adapter to adapt the pre-trained model to new tasks, which is orthogonal to the prompt-based methods. Thus we do not include their results in our comparison.

4.4 Comparison with Prompt-DT type of methods

Table 1: Results for MuJoCo control tasks and MW tasks. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. The dataset we utilized is the full dataset. We test all the results on unseen tasks with three random seeds. LPDT outperforms baselines on the Cheetah environment and is competitive in the Ant environment.

Task	Prompt-DT	Prompt-Tuning DT	Prompt Diffuser	LPDT-Classifier	LPDT-InfoNCE
Cheetah-dir	933.91±7.04	941.5±3.2	945.3±7.2	947.84±1.53	951.72±4.08
Cheetah-vel	-34.71±2.80	-40.1±3.8	-35.3±2.4	-31.57±2.70	-35.98±7.15
Ant-dir	396.07±9.78	427.9±4.3	432.1±6.7	374.13±23.05	412.47±21.01
MW reach-v2	692.29±9.32	472.5±29.0	555.7±6.8	497.61 ± 48.15	528.21±114.24
MW pick-place-v2	3773.82± 356.05	-	-	3508.12±243.93	3543.38±191.32

In this section, we conduct experiments on our proposed LPDT and baseline methods to evaluate their performance. In addition, we compare variants of LPDT with different prompt regularization. The average episode accumulated reward in the test task set serves as the metric for all methods. We directly compare our approaches, including the supervised classifier version and the unsupervised InfoNCE version, with prior works. We test the model with three different seeds and record the

average return over all the testing tasks. Since there is no published code for Prompt-Tuning DT and Prompt Diffuser, we report the results from their papers for a fair comparison.

The results are summarized in Table 1. Note that Prompt-Tuning DT and Prompt Diffuser are two prompt-tuning methods with leading performance among existing Prompt-based works. Table 1 demonstrates that our LPDT outperforms the baseline algorithms on MuJoCo Control tasks but suffers from inferior performance in Meta World compared with Prompt-DT. The possible limitation may be due to the huge difference between the RL task and the language task. Table 1 also illustrates that our approach performs better than the baselines in Cheetah-dir and Cheetah-vel, while it is not as good as Prompt Diffuser in Ant-dir but still better than Prompt-DT. For the results in the Meta World task, we report the results of Prompt-Tuning DT and Prompt Diffuser from their papers.




Figure 2: Results on MuJoCo controls with the Cheetah-dir, Cheetah-vel and Ant-dir for Prompt-DT and our two methods LPDT-Classifier and LPDT-InforNCE. The dataset we utilized is the full dataset. We plot the figures on unseen tasks with the average returns with one seed and 20 evaluation episodes. The figure demonstrates that our LPDT needs less sample data compared with Prompt-DT to achieve superior performance.

Figure 2 illustrates the evaluation of Prompt-DT and our two LPDT approaches. All the plotted methods are tested through 50 episode returns on the unseen tasks. The tasks Cheetah-dir, Cheetah-vel, and Ant-dir have prompts of length $K^* = 5$. Figure 2 shows that our approaches can outperform the baseline method Prompt-DT and demonstrate that they need fewer sample data compared with Prompt-DT to achieve superior performance.

5 Conclusion

In this work, we proposed a novel framework for improving the few-shot prompt ability of decision transformers in offline reinforcement learning, i.e., Language model-initialized Prompt Decision Transformer (LPDT). By leveraging pre-trained language models and combining them with domain-specific RL datasets, LPDT demonstrates improved few-shot prompt capabilities and outperforms or is competitive with best existing baselines in prompt based methods in terms of cumulative rewards on unseen tasks. Our approach has the potential to significantly reduce the data requirements for offline RL tasks, making it more applicable to real-world scenarios where collecting large amounts of RL trajectories is challenging. Furthermore, our results highlight the importance of using pre-trained language models as a starting point for decision-making tasks and demonstrate the effectiveness of our prompt regularization methods in enhancing the model’s ability to distinguish task information contained in prompts.

While LPDT has shown promising results, there are several limitations to our approach. Due to computing resource constraints, our language models are currently limited to GPT-2 and DistilGPT-2. To fully realize the potential of pre-trained language models for decision-making tasks, we hope to extend our approach to more open-source language models and utilize more efficient fine-tuning techniques in the future. Additionally, exploring alternative architectures or incorporating multi-task learning could further enhance the performance of LPDT. Future work will focus on addressing these limitations and expanding the scope of LPDT to a broader range of pre-trained language models and decision-making tasks. As such, LPDT offers a promising direction for future research in offline RL.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
- Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, pp. 2052–2062. PMLR, 2019.
- Hiroki Furuta, Yutaka Matsuo, and Shixiang Shane Gu. Generalized decision transformer for offline hindsight information matching. *arXiv preprint arXiv:2111.10364*, 2021.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Shengchao Hu, Li Shen, Ya Zhang, and Dacheng Tao. Prompt-tuning decision transformer with preference ranking. *arXiv preprint arXiv:2305.09648*, 2023.
- Shengchao Hu, Li Shen, Ya Zhang, and Dacheng Tao. Prompt tuning with diffusion for few-shot pre-trained policy generalization, 2024. URL <https://openreview.net/forum?id=7rex81EZH2>.
- Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *Advances in neural information processing systems*, 32, 2019.
- Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.
- Aviral Kumar, Anikait Singh, Stephen Tian, Chelsea Finn, and Sergey Levine. A workflow for offline model-free robotic reinforcement learning. *arXiv preprint arXiv:2109.10813*, 2021.
- Kuang-Huei Lee, Ofir Nachum, Mengjiao Sherry Yang, Lisa Lee, Daniel Freeman, Sergio Guadarrama, Ian Fischer, Winnie Xu, Eric Jang, Henryk Michalewski, et al. Multi-game decision transformers. *Advances in Neural Information Processing Systems*, 35:27921–27936, 2022.

- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, et al. Pre-trained language models for interactive decision-making. *Advances in Neural Information Processing Systems*, 35:31199–31212, 2022.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- Tianwei Ni, Michel Ma, Benjamin Eysenbach, and Pierre-Luc Bacon. When do transformers shine in rl? decoupling memory from credit assignment. *Advances in Neural Information Processing Systems*, 36, 2024.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Raul Puri and Bryan Catanzaro. Zero-shot text classification with generative language models. *arXiv preprint arXiv:1912.10165*, 2019.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Machel Reid, Yutaro Yamada, and Shixiang Shane Gu. Can wikipedia help offline reinforcement learning? *arXiv preprint arXiv:2201.12122*, 2022.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- Ruizhe Shi, Yuyao Liu, Yanjie Ze, Simon S Du, and Huazhe Xu. Unleashing the power of pre-trained language models for offline reinforcement learning. *arXiv preprint arXiv:2310.20587*, 2023.
- Samarth Sinha, Ajay Mandlekar, and Animesh Garg. S4rl: Surprisingly simple self-supervision for offline reinforcement learning in robotics. In *Conference on Robot Learning*, pp. 907–917. PMLR, 2022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Zhihui Xie, Zichuan Lin, Deheng Ye, Qiang Fu, Yang Wei, and Shuai Li. Future-conditioned unsupervised pretraining for decision transformer. In *International Conference on Machine Learning*, pp. 38187–38203. PMLR, 2023.
- Mengdi Xu, Yikang Shen, Shun Zhang, Yuchen Lu, Ding Zhao, Joshua Tenenbaum, and Chuang Gan. Prompting decision transformer for few-shot policy generalization. In *international conference on machine learning*, pp. 24631–24645. PMLR, 2022.
- Mengdi Xu, Yuchen Lu, Yikang Shen, Shun Zhang, Ding Zhao, and Chuang Gan. Hyper-decision transformer for efficient online policy adaptation. *arXiv preprint arXiv:2304.08487*, 2023a.
- Mengdi Xu, Yuchen Lu, Yikang Shen, Shun Zhang, Ding Zhao, and Chuang Gan. Hyper-decision transformer for efficient online policy adaptation. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=AatUEvC-Wjv>.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, pp. 1094–1100. PMLR, 2020.

Xiangyuan Zhang, Weichao Mao, Haoran Qiu, and Tamer Başar. Decision transformer as a foundation model for partially observable continuous control. *arXiv preprint arXiv:2404.02407*, 2024.

Qinqing Zheng, Amy Zhang, and Aditya Grover. Online decision transformer. In *international conference on machine learning*, pp. 27042–27059. PMLR, 2022.

A Related Work

Decision Transformer. Decision Transformer (DT) (Chen et al., 2021) emerged as a type of algorithm for offline RL by using the powerful Transformer architecture for decision-making. DT models RL trajectories as a sequence generation problem and utilizes the next-token generation paradigm for training. Thus, DT takes the history tokens such as the return, state, and action to predict the next action, which formulates the decision-making as an action prediction or sequence generation in a supervised fashion. Since DT can fully utilize the whole trajectories and is easy to train compared with dynamic programming-based offline RL, many of the following works improved the performance under different settings. For example, Lee et al. (2022) proposed the Multi-game Decision Transformer which is trained on part of the Atari games as the multi-tasks training and fine-tuned on the remaining games to achieve efficient adaptation. Hyper Decision Transformer (Xu et al., 2023a) adds an adapter into Decision Transformer and is fine-tuned on unseen tasks through the demonstration without expert actions. Xie et al. (2023) proposed to predict the action conditioned on the future trajectory embedding instead of conditioned on the return. Trajectory Transformer (Janner et al., 2021) is another research line, which is trained on sequences of state, action, and rewards and generated with the beam search.

Prompt-DT. Prompt Decision Transformer (Xu et al., 2022) utilizes the prompt-based framework to do the meta-RL. It is trained on multi-RL tasks with offline datasets. During the training, the prompts or demonstrations which are a small part of the trajectory are combined with trajectories. During the testing on unseen tasks, the prompt can be a guide for indicating the tasks and help the model predict the action to interact with the environments. Following Prompt-DT, several works are adopting the prompt tuning method to achieve a high-quality prompt. Prompt-Tuning DT (Hu et al., 2021) uses the preference ranking function and black-box tuning method to tune the prompt when testing on unseen tasks to achieve a high-quality prompt. Moreover, Prompt Diffuser (Hu et al., 2024) leverages the diffusion model to generate high-quality prompts leading to improved performance in downstream RL tasks. Different from these works, we adopt the prompt regularization which aims to learn a high-quality prompt embedding to distinguish the different but similar RL tasks. Our method adopts this regularization during the training procedure in the prompt dataset.

Language model based DT. Large language models have achieved many surprising effects in various tasks in recent years. Pre-trained on large datasets such as the corpus of the Internet, LLMs such as GPTs (Radford et al., 2019) demonstrate prompt ability which can generate the text with the guide of the task information. The success of the large language models motivates the increasing use of pre-trained language models in improving Decision Transformer to solve RL tasks (Chen et al., 2021). Several works utilize the powerful representation generalization ability of language models as policies to do the decision-making. Li et al. (2022) proposed to adopt the pre-trained language models for interactive decision-making to convert the policies to sequence data. Wik-RL (Reid et al., 2022) uses a pre-trained language model from the next-token generation paradigm as the initialization of DT for offline RL tasks. However, it suffers from inferior performance than directly using DT. To overcome these challenges and unleash the power of language models, Shi et al. (2023) proposed the LaMo algorithm which uses a pre-trained language model and parameter-efficient fine-tuning methods to improve the original DT. Zhang et al. (2024) also proposed to use LaMo in partially observable continuous control problems which demonstrates a strong generalization ability. Unlike all the above methods, our approach is fine-tuned for learning to identify different prompts for various RL tasks. And during the testing phase, just a small part of the trajectories is used in our method as the prompt without updating the model.

B Details on the Experiment Environments

We evaluate our approach on the MuJoCo tasks and Meta-World ML1 tasks. We split the tasks in these environments into the training sets and the testing sets. The tasks in Cheetah-dir and Ant-dir are split by directions. The tasks in Cheetah-vel are split by the goal velocities. In Meta-World, the tasks are defined by different goal positions. The detailed task indexes can be found in Table 2. The experiments we conducted all followed this setting which guaranteed consistency during the evaluation.

Table 2: Training and testing task indexes when testing the generalization ability. We follow the tasks split between Prompt-DT and previous works to guarantee a direct comparison with baselines.

Environment	Number of tasks	Tasks indexes
Cheetah-dir	Training set: 2	[0,1]
	Testing set: 2	[0,1]
Cheetah-vel	Training set: 35	[0-1,3-6,8-14,16-22,24-25,27-39{}}]
	Testing set: 5	[2,7,15,23,26]
Ant-dir	Training set: 45	[0-5,7-16,18-22,24-29,31-40,42-49]
	Testing set: 5	[6,17,23,30,41]
Meta-World reach-v2	Training set: 15	[1-5,7,8,10-14,17-19]
	Testing set: 5	[6,9,15,16,20]
Meta-World pick-place-v2	Training set: 15	[0-10, 12-16,28-24,25-35,37-41,41-40]
	Testing set: 5	[11,17,25,36,41]

C Hyperparameters

In this section, we show the hyperparameter of our LPDT algorithm for experiments presented in Table 1. The hyperparameters have two parts, corresponding to the Transformer architecture and the prompt regularization respectively. We list these hyperparameters in Table 3.

Table 3: Detail on hyperparameters used in our experiments in Table 1. We show that the hyperparameters are in two parts which are parameters for the model backbone and prompt regularization respectively.

Hyperparameters	Value
K (length of context τ)	20
Training batch size for each task	16
Number of evaluation episodes for each task	20
Learning rate	1e-4
Learning rate decay weight	1e-4
Language initialization	DistilGPT2
Embedding dimension	128
Activation	ReLU
Classifier hyperparameter	0.1
Classifier layers	2
Classifier MLP dimension	128
InfoNCE hyperparameter	0.1
InfoNCE temperature	1
InfoNCE MLP dimension	128

D More Experimental Results

In this section, we provide comprehensive summaries of the experimental results. These are the results of all the unseen tasks in MuJoCo control environments. This summary includes all the experiments and ablation study results on our various components. We show that our methods with prompt regularization are much better than those without regularization and with text regularization. Table 4 demonstrates the results in Cheetah-dir, Table 5 refers to the results in Cheetah-vel and Table 6 refers to the Ant-dir. These results further support our observations and conclusions drawn in the experiment section.

Table 4: Results on Cheetah-dir. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. The dataset we utilized is the full dataset. We test all the results on unseen tasks with three random seeds. Notably, our approach LPDT outperforms the baselines on the Cheetah-dir environment.

Methods	Cheetah-dir-0	Cheetah-dir-1	Average
Prompt-DT	669.79±5.68	1198.03±19.12	933.91±7.04
LPDT w/o regularization	686.19±3.30	1202.41±14.83	944.30±6.31
LPDT-Text	686.58±2.76	1201.61±1.57	944.09±1.68
LPDT-Classfer	692.63±3.57	1203.04±6.40	947.84±1.53
LPDT-InfoNCE	690.66±3.85	1212.78±5.26	951.72±4.08

Table 5: Results on the Cheetah-vel. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. The dataset we utilized is the full dataset. We test all the results on unseen tasks with three random seeds. Notably, our approach LPDT outperforms the baselines on the Cheetah-vel environment.

Methods	Cheetah-vel-2	Cheetah-vel-7	Cheetah-vel-15	Cheetah-vel-23	Cheetah-vel-26	Average
Prompt-DT	-54.48±1.76	-18.12±7.84	-35.92±6.22	-25.93±0.21	-39.11±5.31	-34.71±2.80
LPDT w/o regularization	-36.04±5.46	-18.06±5.40	-39.31±11.77	-51.87±1.55	-35.71±7.21	-36.20±4.05
LPDT-Text	-34.28±3.17	-18.01±11.37	-36.69±12.79	-53.91±3.45	-64.92±31.36	-41.56±9.65
LPDT-Classfer	-35.26±9.09	-13.86±5.02	-24.21±8.99	-48.05±13.96	-36.48±5.44	-31.57±2.70
LPDT-InfoNCE	-38.54±11.37	-13.11±4.30	-33.58±8.23	-37.79±2.19	-56.89±23.33	-35.98±7.15

Table 6: Results on Ant-dir. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. The dataset we utilized is the full dataset. We test all the results on unseen tasks with three random seeds. Notably, our approach LPDT outperforms the baselines on the Ant-dir environment.

Methods	Ant-dir-6	Ant-dir-17	Ant-dir-23	Ant-dir-30	Ant-dir-41	Average
Prompt-DT	628.22±119.58	413.76±4.50	361.18±49.51	358.81±2.30	384.37±20.77	396.07±9.78
LPDT w/o regularization	362.08±64.18	411.79±6.13	324.34±98.67	381.20±1.79	325.36±1.84	360.95±11.07
LPDT-Text	254.53±18.46	418.98±8.00	362.91±54.60	371.94±4.79	348.19±46.48	351.31±23.13
LPDT-Classfer	398.80±120.62	417.85± 4.74	342.60± 82.08	365.02± 1.14	346.37±26.96	374.13± 23.04
LPDT-InfoNCE	555.18±164.12	418.42±5.43	376.34±15.52	362.91±5.22	349.52±46.83	412.47±21.01

E Ablation Studies

In this section, we provide ablation studies on LPDT to test the role of prompt regularization and language initialization respectively.

Table 7: Results for MuJoCo control tasks and MW tasks with different regularization methods of our method including w/o regularization, text regularization, classifier regularization and InfoNCE regularization. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. We test all the results on unseen tasks with three random seeds. The dataset we utilized is the full dataset. We demonstrate that the regularization on prompt can help distinguish the task and improve the performance compared with the method without regularization.

Task	Prompt-DT	LPDT w/o regularization	LPDT-Text	LPDT-classifier	LPDT-InfoNCE
Cheetah-dir	933.91±7.04	944.30±6.31	944.09±1.68	947.84±1.53	951.72±4.08
Cheetah-vel	-34.71±2.80	-36.20±4.05	-41.56±9.65	-31.57±2.70	-35.98±7.15
Ant-dir	396.07±9.78	360.95±11.07	351.31±23.13	374.13±23.05	412.47±21.01
MW reach-v2	692.29±9.32	431.87±115.19	459.65±76.01	497.61±48.15	528.21±114.24
MW pick-place-v2	3773.82± 356.05	3700.34±68.45	3678.53±58	3508.12±243.93	3543.38±191.32

E.1 The role of prompt regularization

We first compare our proposed model with the same approach without the prompt regularization, denoted as LPDT w/o regularization. We also follow the settings of LaMo (Shi et al., 2023) and add a text regularization introduced by LaMo for comparison. We summarize the results in Table 7. Specifically, we find that prompt regularization helps the model distinguish between tasks and improve performance. LPDT without regularization or with text regularization suffers from inferior performance compared to our prompt regularization methods. Note that text regularization is trained with an NLP dataset, which is time-consuming and inefficient.

E.2 Data efficiency of language initialization

To verify whether language initialization can incorporate prior knowledge about the downstream RL tasks, we split the dataset to train Prompt-DT and our approaches. We also evaluate the results on a portion of the dataset using a ratio of 0.1. The results are summarized in Table 8.

Table 8: Ratio results for MuJoCo control tasks and MW tasks with the full dataset and 0.1 ratio (10%) dataset. The best mean scores are highlighted in bold. For each environment, the length of the prompt is $K^* = 5$. We test all the results on unseen tasks with three random seeds. We demonstrate that language initialization can improve the performance of our LPDT.

Task	Datset	Cheetah-dir	Cheetah-vel	Ant-dir	MW reach-v2	MW pick-place-v2
Prompt-DT	Full	933.91±7.04	-34.71±2.80	396.07±9.78	692.29±9.32	3773.82±356.05
	0.1	890.35±11.95	-42.46±9.25	361.13±4.46	601.56±40.44	3062.67±283.56
LPDT w/o regularization	Full	944.30±6.31	-36.20±4.05	360.95±11.07	431.87±115.19	3700.34±68.45
	0.1	927.80±4.91	-37.25±5.85	353.56±27.98	446.03±30.94	3555.53±255.52
LPDT-Classifier	Full	947.84±1.53	-31.57±2.70	374.13±23.05	497.61±48.15	3508.12±243.93
	0.1	931.42±6.40	-34.37±8.58	347.12±28.48	457.94±71.38	3118.57±77.47
LPDT-InfoNCE	Full	951.72±4.08	-35.98±7.15	412.47±21.01	528.21±114.24	3543.38±191.32
	0.1	919.75±4.17	-34.20±8.21	369.73±21.97	449.04±44.13	3341.35±131.93

Table 8 illustrates the results under different sizes of training datasets. We adopt DistilGPT2 for language initialization without any regularization. We also use LoRA to fine-tune the language models to the MuJoCo dataset. The results show that with only a 0.1 ratio (10%) of the dataset, the methods with language initialization outperform Prompt-DT when tested on unseen tasks, demonstrating that the language pre-trained model contains valuable knowledge for downstream RL tasks.

Combining these two ablation studies, we conclude that language initialization can improve performance when data is limited. Furthermore, the prompt regularization method can help the language model perform better when tested on unseen RL tasks.