
Online Decision Transformer

Qinqing Zheng¹ Amy Zhang^{1,2} Aditya Grover^{1,3}

Abstract

Recent work has shown that offline reinforcement learning (RL) can be formulated as a sequence modeling problem (Chen et al., 2021; Janner et al., 2021) and solved via approaches similar to large-scale language modeling. However, any practical instantiation of RL also involves an online component, where policies pretrained on passive offline datasets are finetuned via task-specific interactions with the environment. We propose Online Decision Transformers (ODT), an RL algorithm based on sequence modeling that blends offline pretraining with online finetuning in a unified framework. Our framework uses sequence-level entropy regularizers in conjunction with autoregressive modeling objectives for sample-efficient exploration and finetuning. Empirically, we show that ODT is competitive with the state-of-the-art in absolute performance on the D4RL benchmark but shows much more significant gains during the finetuning procedure.

1 Introduction

Generative pretraining for sequence modeling has emerged as a unifying paradigm for machine learning in a number of domains and modalities, notably in language and vision (Radford et al., 2018; Chen et al., 2020; Brown et al., 2020; Lu et al., 2022). Recently, such a pretraining paradigm has been extended to offline reinforcement learning (RL) (Chen et al., 2021; Janner et al., 2021), wherein an agent is trained to autoregressively maximize the likelihood of trajectories in the offline dataset. During training, this paradigm essentially converts offline RL to a supervised learning problem (Schmidhuber, 2019; Srivastava et al., 2019; Emmons et al., 2021). However, these works present an incomplete picture as policies learned via offline RL are limited by the quality of the training dataset and need to be

finetuned to the task of interest via online interactions. It remains an open question whether such supervised learning paradigm can be extended to online settings.

Unlike language and perception, online finetuning for RL is fundamentally different from the pretraining phase as it involves data acquisition via *exploration*. The need for exploration renders traditional supervised learning objectives (e.g., mean squared error) for offline RL insufficient in the online setting. Moreover, it has been observed that for standard online algorithms, access to offline data can often have zero or even negative effect on the online performance (Nair et al., 2020). Hence, the overall pipeline for offline pretraining followed by online finetuning for RL policies needs a careful consideration of training objectives and protocols.

We introduce Online Decision Transformers (ODT), a learning framework for RL that blends offline pretraining with online finetuning for sample-efficient policy optimization. Our framework builds on the decision transformer (DT) (Chen et al., 2021) architecture previously introduced for offline RL and is especially catered to scenarios where online interactions can be expensive which necessitates both offline pretraining and sample-efficient finetuning. We identify several key shortcomings with DTs that are incompatible with online learning and rectify them, leading to superior performance for our overall pipeline.

First, we shift from deterministic to stochastic policies for defining exploration objectives during the online phase. We quantify exploration via the entropy of the policy similar to max-ent RL frameworks (Levine, 2018). Unlike traditional frameworks however, the policy entropy for ODT is constrained at an aggregate level over trajectories (as opposed to individual time-steps) and its dual form regularizes a supervised learning objective (as opposed to direct return maximization). Next, we develop a novel replay buffer (Mnih et al., 2015) consistent with the architecture and training protocol of ODT. The buffer stores trajectories and is populated via online rollouts from ODT. Since ODT parameterizes return-conditioned policies, we further investigate strategies for specifying the desired returns during online rollouts. This value however might not match with the true returns observed during a rollout. To address this challenge, we extend a notion of hindsight experience replay (Andrychowicz et al., 2017) to our setting and relabel rolled out trajectories with the corrected return tokens

¹Meta AI Research ²University of California, Berkeley
³University of California, Los Angeles. Correspondence to: Qinqing Zheng <zhengqinqing@gmail.com>.

before augmenting them.

Empirically, we validate our overall framework by comparing its performance with state-of-the-art algorithms on the D4RL benchmark (Fu et al., 2020). We find that our relative improvements due to our finetuning strategy outperform other baselines (Nair et al., 2020; Kostrikov et al., 2021b), while exhibiting competitive absolute performance when accounting for the pretraining results of the base model. Finally, we supplement our main results with rigorous ablations and additional experimental designs to justify and validate the key components of our approach.

2 Related Work

Our work encompasses two broad avenues of research which we detail here.

Transformers for RL. There has been much recent exciting progress formulating the offline RL problem as a context-conditioned sequence modeling problem (Chen et al., 2021; Janner et al., 2021). These works build on the reinforcement learning as supervised learning paradigm (Schmidhuber, 2019; Srivastava et al., 2019; Emmons et al., 2021) that focuses on predictive modeling of action sequences conditioned on a task specification (e.g., target goal or returns) as opposed to explicitly learning Q-functions or policy gradients. Chen et al. (2021) trains a transformer (Vaswani et al., 2017; Radford et al., 2018) as a model-free context-conditioned policy, and Janner et al. (2021) trains a transformer as both a policy and model and show that beam search can be used to improve upon purely model-free performance. However, these works only explore the offline RL setting, which is similar to the fixed datasets that transformers are traditionally trained with in natural language processing applications. Our work focuses on extending these results to the online finetuning setting, showing competitiveness with state-of-the-art RL methods.

Offline RL methods primarily add a conservative component to an existing off-policy RL method to prevent out-of-distribution extrapolation, but require many tweaks and re-tuning of hyperparameters to work (Kumar et al., 2020a; Kostrikov et al., 2021a). Similar to our work, Fujimoto & Gu (2021) show the benefits of adding a behavior cloning term to offline RL methods, and that the simple addition of this term allows the porting of off-policy RL algorithms to the offline setting with minimal changes.

Offline RL with Online Finetuning. While ODT stems from a different perspective than traditional RL methods, there is much existing work focused on the same paradigm of pre-training on a given offline dataset and finetuning in an online environment. Nair et al. (2020) showed that naïve application of offline or off-policy RL methods to the offline pre-training and online finetuning regime often does not help, or even hinders, performance. This poor performance

in off-policy methods can be attributed to off-policy bootstrapping error accumulation (Munos, 2003; 2005; Farahmand et al., 2010; Kumar et al., 2019). In offline RL methods, poor performance in the online finetuning regime can be explained by excess conservatism, which is necessary in the offline regime to prevent value overestimation of out-of-distribution states. Nair et al. (2020) was the first to propose an algorithm that works well for both the offline and online training regimes. Recent work (Kostrikov et al., 2021b) also proposes an expectile-based implicit Q-learning algorithm for offline RL that also shows strong online finetuning performance, because the policy is extracted via a behavior cloning step that avoids out-of-distribution actions.

Lee et al. (2021) tackles the offline-online setting with a balanced replay scheme and an ensemble of Q functions to maintain conservatism during offline training. Lu et al. (2021) improves upon AWAC (Nair et al., 2020), which exhibits collapse during the online finetuning stage, by incorporating positive sampling and exploration during the online stage. We also find that positive sampling and exploration are key to good online finetuning, but we will show how these traits naturally occur in ODT, leading to a simple, end-to-end method that automatically adapts to both offline and online settings.

3 Preliminaries

We assume our environment can be modeled as a Markov decision process (MDP), which can be described as $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, p, P, R, \gamma \rangle$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $P(s_{t+1}|s_t, a_t)$ is the probability distribution over transitions, $R(s_t, a_t)$ is the reward function, and γ is the discount factor (Bellman, 1957). An agent starts in a initial state s_1 sampled from a fixed distribution $p(s_1)$, then at each timestep t it takes an action $a_t \in \mathcal{A}$ from a state $s_t \in \mathcal{S}$ and moves to a next state $s_{t+1} \sim P(\cdot|s_t, a_t)$. After each action the agent receives a deterministic reward $r_t = R(s_t, a_t)$. Note that our algorithms also directly apply to partially observable Markov decision processes (POMDP), but we use the MDP framework for ease of exposition.

3.1 Setup and Notation

We are interested in online finetuning of Decision Transformer (DT) (Chen et al., 2020), wherein an agent will have access to a non-stationary training data distribution \mathcal{T} . Initially, during pretraining, \mathcal{T} corresponds to the offline data distribution $\mathcal{T}_{\text{offline}}$ and is accessed via an offline dataset $\mathcal{J}_{\text{offline}}$. During finetuning it is accessed via a replay buffer $\mathcal{J}_{\text{replay}}$. Let τ denote a trajectory and let $|\tau|$ denote its length. The return-to-go (RTG) of a trajectory τ at timestep t , $g_t = \sum_{t'=t}^{|\tau|} r_{t'}$, is the sum of future rewards from that timestep. Let $\mathbf{a} = (a_1, \dots, a_{|\tau|})$, $\mathbf{s} = (s_1, \dots, s_{|\tau|})$ and $\mathbf{g} = (g_1, \dots, g_{|\tau|})$ denote the sequence of action, states and RTGs of τ , respectively.

3.2 Decision Transformer

Decision Transformer processes a trajectory τ as a sequence of 3 types of input tokens: RTGs, states and actions: $(g_1, s_1, a_1, g_2, s_2, a_2, \dots, g_{|\tau|}, s_{|\tau|}, a_{|\tau|})$. Specifically, the initial RTG g_1 is equal to the return of the trajectory. At timestep t , DT uses the tokens from the latest K timesteps to generate an action a_t . Here, K is a hyperparameter and is also referred to as the *context length* for the transformer. Note that the context length during evaluation can be shorter than the one used for training, as we will demonstrate later in our experiments. DT learns a deterministic policy $\pi_{\text{DT}}(a_t | \mathbf{s}_{-K,t}, \mathbf{g}_{-K,t})$, where $\mathbf{s}_{-K,t}$ is shorthand for the sequence of K past states $s_{\max(1,t-K+1):t}$ and similarly for $\mathbf{g}_{-K,t}$. This is an autoregressive model of order K . In particular, DT parameterized the policy through a GPT architecture (Radford et al., 2018), which applies a causal mask to enforce such autoregressive structure in the predicted action sequence.

For simplicity, we assume the data distribution \mathcal{T} generates associated length- K action, state and RTG subsequences, which are from the same trajectory. With a little abuse of notation, we also use $(\mathbf{a}, \mathbf{s}, \mathbf{g})$ to denote the sample from \mathcal{T} . This allows us to easily present the training objective of our approach, and the above notations readily apply here. The policy is trained to predict the action tokens under the standard ℓ_2 loss

$$\mathbb{E}_{(\mathbf{a}, \mathbf{s}, \mathbf{g}) \sim \mathcal{T}} \left[\frac{1}{K} \sum_{k=1}^K (a_k - \pi_{\text{DT}}(\mathbf{s}_{-K,k}, \mathbf{g}_{-K,k}))^2 \right]. \quad (1)$$

In practice, we uniformly sample the length- K subsequences from the offline dataset $\mathcal{J}_{\text{offline}}$ (or the replay buffer $\mathcal{J}_{\text{replay}}$ during finetuning), see Appendix F.

For evaluation, we specify the desired performance g_1 and an initial state s_1 . DT then generates the action $a_1 = \pi_{\text{DT}}(s_1, g_1)$. Once an action a_t is generated, we execute it and observe the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$ and obtain a reward $r_t = R(s_t, a_t)$. This gives us the next RTG as $g_{t+1} = g_t - r_t$. As before, DT generates a_2 based on s_1, s_2 and g_1, g_2 . This process is repeated until the episode terminates.

4 Online Decision Transformer

RL policies trained on purely offline datasets are typically sub-optimal due to the limitations of training data as the offline trajectories might not have high return and cover only a limited part of the state space. One natural strategy to improve performance is to finetune the pretrained RL agents via online interactions. However, the learning formulation for a standard decision transformer is insufficient for online learning, and as we shall show in our experiment ablations, collapses when used naively for online data acquisition. In this section, we introduce key modifications to decision transformers for enabling sample-efficient online finetuning.

As a first step, we present a generalized, probabilistic learning objective. We will extend this formulation to account for exploration in online decision transformers (ODT). In the probabilistic setup, our goal is to learn a *stochastic* policy that maximizes the likelihood of the dataset. For example, for continuous action spaces, we can use the standard choice (Haarnoja et al., 2018a; Fujimoto & Gu, 2021; Kumar et al., 2020b; Emmons et al., 2021) of a multivariate Gaussian distribution with a diagonal covariance matrix to model the action distributions conditioned on states and RTGs. Let θ denote the policy parameters. Formally, our policy is

$$\begin{aligned} & \pi_{\theta}(a_t | \mathbf{s}_{-K,t}, \mathbf{g}_{-K,t}) \\ &= \mathcal{N}(\mu_{\theta}(\mathbf{s}_{-K,t}, \mathbf{g}_{-K,t}), \Sigma_{\theta}(\mathbf{s}_{-K,t}, \mathbf{g}_{-K,t})), \quad \forall t, \end{aligned} \quad (2)$$

where the covariance matrix Σ_{θ} is assumed to be diagonal. Given a stochastic policy, we maximize the log-likelihood of the trajectories in the training dataset, or equivalently minimize the negative log-likelihood (NLL) loss¹

$$\begin{aligned} J(\theta) &= \frac{1}{K} \mathbb{E}_{(\mathbf{a}, \mathbf{s}, \mathbf{g}) \sim \mathcal{T}} [-\log \pi_{\theta}(\mathbf{a} | \mathbf{s}, \mathbf{g})] \\ &= \frac{1}{K} \mathbb{E}_{(\mathbf{a}, \mathbf{s}, \mathbf{g}) \sim \mathcal{T}} \left[-\sum_{k=1}^K \log \pi_{\theta}(a_k | \mathbf{s}_{-K,k}, \mathbf{g}_{-K,k}) \right]. \end{aligned} \quad (3)$$

The policies we consider here subsume the deterministic policies considered by DT. Optimizing the objective (1) is equivalent to optimizing (3) assuming the covariance matrix Σ_{θ} is diagonal and the variances are the same across all the dimensions, which is a special case covered by our assumption.

4.1 Max-Entropy Sequence Modeling

The key property of an online RL algorithm is the ability to balance the exploration-exploitation trade-off. Even with stochastic policies, traditional formulation of a DT as in Eq. (3) does not account for exploration. To address this shortcoming, we first quantify exploration via the policy entropy¹ defined as:

$$\begin{aligned} H_{\theta}^{\mathcal{T}}[\mathbf{a} | \mathbf{s}, \mathbf{g}] &= \frac{1}{K} \mathbb{E}_{(\mathbf{s}, \mathbf{g}) \sim \mathcal{T}} [H[\pi_{\theta}(\mathbf{a} | \mathbf{s}, \mathbf{g})]] \\ &= \frac{1}{K} \mathbb{E}_{(\mathbf{s}, \mathbf{g}) \sim \mathcal{T}} \left[\sum_{k=1}^K H[\pi_{\theta}(a_k | \mathbf{s}_{-K,k}, \mathbf{g}_{-K,k})] \right], \end{aligned} \quad (4)$$

where $H[\pi_{\theta}(a_k)]$ denotes the Shannon entropy of the distribution $\pi_{\theta}(a_k)$. The policy entropy depends on the data distribution \mathcal{T} , which is static in the offline pretraining phase but dynamic during finetuning as it depends on the online data acquired during exploration.

Similar to many existing max-ent RL algorithms (Levine, 2018) such as Soft Actor Critic (SAC, Haarnoja et al. (2018a;b)), we explicitly impose a lower bound on the policy entropy to encourage exploration. That is, we are interested in solving the following constrained problem:

$$\min_{\theta} J(\theta) \quad \text{subject to} \quad H_{\theta}^{\mathcal{T}}[\mathbf{a} | \mathbf{s}, \mathbf{g}] \geq \beta, \quad (5)$$

¹We scale both $J(\theta)$ and $H_{\theta}^{\mathcal{T}}[\mathbf{a} | \mathbf{s}, \mathbf{g}]$ by $\frac{1}{K}$. As we discuss later, this allows us to easily compare our objective with SAC.

where β is a prefixed hyperparameter. Following Haarnoja et al. (2018b), in practice, we solve the dual problem of (5) to avoid explicitly dealing with the inequality constraint. Namely, we consider the Lagrangian $L(\theta, \lambda) = J(\theta) + \lambda(\beta - H_\theta^T[\mathbf{a}|s, \mathbf{g}])$ and solve the problem $\max_{\lambda \geq 0} \min_{\theta} L(\theta, \lambda)$ by alternately optimizing θ and λ . Optimizing θ with fixed λ is equivalent to

$$\min_{\theta} J(\theta) - \lambda H_\theta^T[\mathbf{a}|s, \mathbf{g}], \quad (6)$$

and optimizing λ with fixed θ boils down to solving

$$\min_{\lambda \geq 0} \lambda(H_\theta^T[\mathbf{a}|s, \mathbf{g}] - \beta). \quad (7)$$

We highlight a few focal points for understanding our method from both the primal and dual point of view.

(1) Problem (6) can be viewed as the corresponding regularized form of our primal problem (5), where the dual variable λ serves the role of temperature variable in many regularized RL formulations. As a key difference to SAC and other classic RL methods, our loss function $J(\theta)$ is the NLL rather than the discounted return. In principal, we focus only on supervised learning of action sequences as opposed to explicitly maximizing returns. Consequently, the objective in Eq. (6) can be interpreted as minimizing the expected difference between the log-probability of the observed actions and λ -scaled log-probability of actions drawn from $\pi_\theta(\cdot|s, \mathbf{g})$. That is, we attempt to learn π_θ so that it matches the observed action distribution with some deviation, and the dual variable λ explicitly controls the degree of mismatch.

(2) Rigorously speaking, $H_\theta^T[\mathbf{a}|s, \mathbf{g}]$ is a cross conditional entropy. This is because the training data distribution \mathcal{T} is generally not the same as the action-state-RTG marginals induced by the current policy π_θ and the transition probability P . During pretraining, \mathcal{T} is the fixed offline data distribution $\mathcal{T}_{\text{offline}}$ whereas during finetuning, as we use a replay buffer to store the collected online data, \mathcal{T} is accessed via a replay buffer that depends on the current policy π_θ and the data gathered by the policies at the previous iterations. However, over the course of training, if the policy converges, this cross entropy term will converge to the entropy. Section 4.3 discusses this in more detail. In other words, our objective function automatically adapts to play a suitable role for both offline and online settings. During offline training, the cross entropy term controls the degree of distribution mismatch whereas during online training, it becomes the entropy and encourages policy exploration.

(3) Another important difference to the classic max-ent RL methods is that our policy entropy $H_\theta^T[\mathbf{a}|s, \mathbf{g}]$, as in Eq. (4), is defined at the level of sequences, as opposed to transitions. Consequently, our constraint in the primal problem (5) differs from the entropy constraint for SAC (Haarnoja et al.,

2018b). For simplicity, let us ignore the RTG variable \mathbf{g} in our policy entropy. While SAC imposes a lower bound β for the policy entropy at all timesteps, our entropy $H_\theta^T[\mathbf{a}|s, \mathbf{g}]$ is averaged over K consecutive timesteps. Hence, our constraint only requires that the entropy averaged over a sequence of K timesteps is lower bounded by β . Therefore, any policy satisfying the transition-level policy constraint of SAC also satisfies our sequence-level constraint. Hence, the space of feasible policies is larger in our case whenever $K > 1$. When $K = 1$, the sequence-level constraint reduces to the transition-level constraint of SAC.

Finally, we make several remarks regarding similarities to SAC with respect to the practical optimization details. First, we do not fully solve the sub-problems (6) and (7). For both of them, we only take one gradient update each time, a.k.a one-step alternating gradient descent. Second, the evaluation of $H_\theta^T[\mathbf{a}|s, \mathbf{g}]$ involves integrals. We approximate each integral using the one-sample Monte Carlo estimation, and the sample is re-parameterized for low-variance gradient computation. As also noted by Haarnoja et al. (2018b), we often observe that the constraint in problem (5) is tight so that the actual entropy $H_\theta^T[\mathbf{a}|s, \mathbf{g}]$ matches β .

4.2 Training Pipeline

We instantiate the formulation described above using a transformer architecture. The online decision transformer (ODT) builds on the DT architecture and incorporates changes due to the stochastic policy. We predict the policy mean and log-variance by two separate fully connected layers at the output. Algorithm 1 summarizes the overall finetuning pipeline in ODT, where the detailed inner training steps are described in Algorithm 2. We outline the major components of these algorithms below and discuss additional design choices and hyperparameters in Appendix B.

Trajectory-Level Replay Buffer. We use a replay buffer to record our past experiences and update it periodically. For most existing RL algorithms, the replay buffer is composed of transitions. After every step of online interaction within a rollout, the policy or the Q-function is updated via a gradient step and the policy is executed to gather a new transition for addition into the replay buffer. For ODT however, our replay buffer consists of trajectories rather than transitions. After offline pretraining, we initialize the replay buffer by the trajectories with highest returns in the offline dataset. Every time we interact with the environment, we completely rollout an episode with the current policy, then refresh the replay buffer using the collected trajectory in a first-in-first-out manner. Afterwards, we update the policy and rollout again, as shown in Algorithm 1. Similar to Haarnoja et al. (2018a), we also observe that evaluating the policy using the mean action generally leads to a higher return, but it is beneficial to use sampled actions for online exploration since it generates more diverse data.

Algorithm 1: Online Decision Transformer

- 1 **Input:** offline data $\mathcal{T}_{\text{offline}}$, rounds R , exploration RTG g_{online} , buffer size N , gradient iterations I , pretrained policy π_{θ}
 - 2 **Initialization:** Replay buffer $\mathcal{T}_{\text{replay}} \leftarrow$ top N trajectories in $\mathcal{T}_{\text{offline}}$.
 - 3 **for** round = 1, ..., R **do**
 - // use randomly sampled actions
 - 4 Trajectory $\tau \leftarrow$ Rollout using \mathcal{M} and $\pi_{\theta}(\cdot|s, \mathbf{g}(g_{\text{online}}))$.
 - 5 $\mathcal{T}_{\text{replay}} \leftarrow \{\mathcal{T}_{\text{replay}} \setminus \{\text{the oldest trajectory}\}\} \cup \{\tau\}$.
 - 6 $\pi_{\theta} \leftarrow$ Finetune ODT on $\mathcal{T}_{\text{replay}}$ for I iterations via Algorithm 2.
-

Algorithm 2: ODT Training

- 1 **Input:** model parameters θ , replay buffer $\mathcal{T}_{\text{replay}}$, training iterations I , context length K , batch size B
 - 2 Compute the trajectory sampling probability $p(\tau) = |\tau| / \sum_{\tau \in \mathcal{T}} |\tau|$.
 - 3 **for** $t = 1, \dots, I$ **do**
 - 4 Sample B trajectories out of $\mathcal{T}_{\text{replay}}$ according to p .
 - 5 **for** each sampled trajectory τ **do**
 - // Hindsight Return Relabeling
 - 6 $\mathbf{g} \leftarrow$ the RTG sequence computed by the true rewards: $\mathbf{g}_t = \sum_{j=t}^{|\tau|} r_j, 1 \leq t \leq |\tau|$.
 - 7 $(\mathbf{a}, \mathbf{s}, \mathbf{g}) \leftarrow$ a length K sub-trajectory uniformly sampled from τ .
 - 8 $\theta \leftarrow$ one gradient update using the sampled $\{(\mathbf{a}, \mathbf{s}, \mathbf{g})\}$.
-

Hindsight Return Relabeling. Hindsight experience replay (HER) is a method for improving the sample-efficiency of goal-conditioned agents in environments with sparse rewards (Andrychowicz et al., 2017; Rauber et al., 2017; Ghosh et al., 2019). The key idea here is to relabel the agent’s trajectories with the achieved goal, as opposed to the intended goal. In the case of ODT, we are learning policies conditioned on an initial RTG. The return achieved during a policy rollout and the induced RTG can however differ from the intended RTG. Inspired by HER, we relabel the RTG token for the rolled out trajectory τ with the achieved returns, such that the RTG token at the last timestep $g_{|\tau|}$ is exactly the reward obtained by the agent $r_{|\tau|}$, see Line 6 of Algorithm 2. This return relabeling strategy applies to environments with both sparse and dense rewards.

RTG Conditioning. ODT requires an hyperparameter, initial RTG g_{online} , for gathering additional online data (see Line 4 of Algorithm 1). Previously, Chen et al. (2021) showed that the actual evaluation return of offline DT has strong correlation with the initial RTG empirically and can often extrapolate to RTG values beyond the maximum returns observed in the offline dataset. For ODT, we find it best to set this hyperparameter to a small, fixed scaling (set to 2 in our experiments) of the expert return. We also experimented with much larger values as well as time-varying curriculum (e.g., quantiles of the best evaluation return in the offline and online datasets) but we found these to be suboptimal relative to a fixed, scaled RTG.

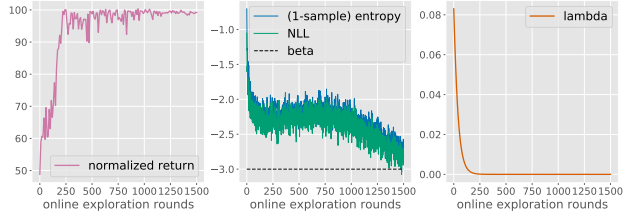


Figure 4.1. An example run of ODT on Hopper task using medium dataset. ODT’s evaluation return converges over the course of training. Meanwhile, both the NLL and the entropy are converging to β , which is -3 in this run. The dual variable λ converges to zero.

Sampling Strategy. Similar to DT, Algorithm 2 uses a two step sampling procedure to ensure that the sub-trajectories of length K in the replay buffer $\mathcal{T}_{\text{replay}}$ are sampled uniformly. We first sample a single trajectory with probability proportional to its length, then uniformly sample a sub-trajectory of length K . For environments with non-negative dense rewards, our sampling strategy is akin to importance sampling. In those cases, the length of a trajectory is highly correlated with its return, as we further highlight in Appendix F.

4.3 Training Dynamics

We comment on some empirical observations regarding the training dynamics of ODT and their implications. We first show an example run where the online return of ODT saturates, suggesting that training has converged. Restricting ourselves to cases where Algorithm 1 converges, we discuss the training dynamics of ODT. This convergence assumption enables us to analyze the simplification of the learning objective over the course of training, and also the change in behavior of the initial RTG token for ODT policies. We emphasize that the convergence guarantee of Algorithm 1 is an open question and beyond the scope of this paper, and our claims will be guided primarily via experiments.

Figure 4.1 plots an example run of ODT. The left panel shows that the return of ODT converges, suggesting a potential convergence of the evaluation policy $\pi_{\theta}(\cdot|s, \mathbf{g}_{\text{eval}})$, where \mathbf{g}_{eval} is the RTG sequence induced by the initial RTG token g_{eval} for evaluation rollouts.

Let us consider the objective function in Eq. (6):

$$\min_{\theta} \overbrace{\mathbb{E}_{(s, \mathbf{a}, \mathbf{g}) \sim \mathcal{T}} [-\log \pi_{\theta}(\mathbf{a}|s, \mathbf{g})]}^{T_{\text{NLL}}} - \lambda \overbrace{\mathbb{E}_{(s, \mathbf{g}) \sim \mathcal{T}} [H(\pi_{\theta}(\cdot|s, \mathbf{g}))]}^{T_{\text{CE}}}. \quad (8)$$

As discussed previously in Section 4.1, the training data distribution \mathcal{T} is static in the offline training stage, but keeps evolving during finetuning. Problem (8) falls into the class of optimization problems where the data distribution is non-stationary and depends on the model parameters. Recent advances in optimization show that the iterates of those problems converge to a fixed point with (stochastic) gradient updates under certain conditions, e.g., see related works including Bottou et al. (2013); Perdomo et al. (2020); Mender-Dünner et al. (2020); Drusvyatskiy & Xiao (2020).

Motivated by our empirical observations and the recent theoretical progress, in the rest of this section, we analyze the training dynamics of ODT under the assumption that the policy learned via Algorithm 1 converges.

To start, we discuss how problem (8) evolves. Its objective contains two terms, T_{NLL} and T_{CE} . As mentioned in Section 4.1, T_{CE} is a cross entropy during pretraining because the expectation is w.r.t. $\mathcal{T} = \mathcal{T}_{\text{offline}}$ (the offline data distribution) rather than the marginal action-state-RTG distribution $\rho_{\pi_\theta}(\mathbf{a}, \mathbf{s}, \mathbf{g})$ induced by π_θ and P . As ODT training converges, a few consequences follow. For simplicity, let us ignore the hindsight return relabeling for now. First, in the online finetuning stage, the training data is sampled from the replay buffer which contains the past exploration rollouts. As the policy π_θ converges, \mathcal{T} will also converge to the policy induced marginal $\rho_{\pi_\theta}(\mathbf{a}, \mathbf{s}, \mathbf{g})$. If this happens, the cross entropy term T_{CE} reduces to the conditional entropy $H_{\rho_{\pi_\theta}}[\mathbf{a}|\mathbf{s}, \mathbf{g}]$, which is also equal to the NLL term T_{NLL} .

As a result, problem (8) reduces to NLL minimization if λ converges to a value between 0 and 1 ($0 \leq \lambda < 1$):

$$\min_{\theta} \mathbb{E}_{(\mathbf{s}, \mathbf{a}, \mathbf{g}) \sim \rho_{\pi_\theta}} [-\log \pi_\theta(\mathbf{a}|\mathbf{s}, \mathbf{g})]. \quad (9)$$

By the complementary slackness (Boyd et al., 2004), the objective of problem (7) converges to zero. For the example showed in Figure 4.1, the constraint is inactive i.e., the entropy is always larger than β , therefore, the dual variable λ converges to zero. We have also observed cases where the constraint is (stochastically) tight, for which λ can be positive. In our experiments, we consistently observe that λ converges to a value between 0 and 1 even when starting from various initial values. In both scenarios, the overall loss will converge to the desired formulation (9). See Appendix G for more discussions.

Put differently, the above objective performs behavior cloning on trajectories generated as per $\rho_{\pi_\theta}(\mathbf{a}, \mathbf{s}, \mathbf{g})$. Now, let us consider hindsight return relabeling and take a closer look at problem (9). Let $\mathbf{g}_{\text{online}}$ denote the RTG subsequence induced by the exploration RTG g_{online} , and \mathbf{g}_{real} denote the real RTG subsequence obtained by relabeling: $\mathbf{g}_{\text{real}, t} = \sum_{i=t}^{|\tau|} r_i$. Problem (9) now becomes

$$\min_{\theta} \mathbb{E}_{\mathbf{s} \sim P, \mathbf{a} \sim \pi_\theta(\cdot|\mathbf{s}, \mathbf{g}_{\text{online}})} [-\log \pi_\theta(\mathbf{a}|\mathbf{s}, \mathbf{g}_{\text{real}}(\mathbf{s}, \mathbf{a}))]. \quad (10)$$

This formulation suggests that we are trying to match the policies $\pi_\theta(\cdot|\mathbf{s}, \mathbf{g}_{\text{real}})$ for all the observed \mathbf{g}_{real} with a single policy $\pi_\theta(\cdot|\mathbf{s}, \mathbf{g}_{\text{online}})$.

Empirically, we make observations in Figure 4.2 that coincide with the above analysis. Here, we vary the initial evaluation RTG g_{eval} for Hopper and inspect the return of ODT after pretraining on `medium` dataset and finetuning, respectively. The return of ODT is strongly correlated with

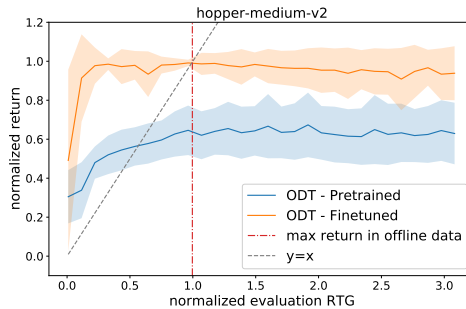


Figure 4.2. The evaluation returns of ODT when conditioned on a wide range of values for g_{eval} .

g_{eval} after the offline pretraining, although the performance saturates after a certain threshold of g_{eval} (around 0.85 in Figure 4.2). In comparison, the finetuned ODT is less sensitive to g_{eval} except for exceptionally large values, where the return slightly declines and the variance increases (cf. Section 5.2). Besides, the threshold for performance saturation gets pushed to an extremely low value (around 0.1 in Figure 4.2). This demonstrates that while the offline pretrained ODT models a relatively wide distribution of policy returns, the distribution learned by ODT after finetuning concentrates on a narrow range of returns, suggesting a potential concentration on a single policy.

5 Experiments

Our experiments aim to answer two primary questions:

1. How does ODT compare with other state-of-the-art approaches for finetuning pretrained policies under a *limited* online budget?
2. How do the individual components of ODT influence its overall performance?

Tasks and Datasets. For answering both these questions, we focus on two types of tasks with offline datasets from the D4RL benchmark (Fu et al., 2020). The first type consists of the Gym locomotion tasks Hopper, Walker, HalfCheetah, and Ant, which are standard environments with dense rewards. For these environments, our evaluation uses the v2 medium and medium-replay datasets which contain trajectories collected by sub-optimal policies. The `medium` dataset contains 1M samples from a policy trained to approximately $\frac{1}{3}$ the performance of an expert policy, and the `medium-replay` dataset uses the replay buffer of a policy trained up to the performance of a medium agent. For the second setup, we consider the goal-reaching tasks AntMaze, where the objective is to move an Ant robot to a target location and the rewards are sparse. The agent obtains reward 1 if it reaches the goal location otherwise 0. We use the v2 umaze and umaze-diverse datasets in our experiments.

5.1 Benchmark Comparisons

For a thorough understanding of different methods, we compare both the offline and online performance. We compare the offline performance of ODT with DT and implicit Q-learning (IQL) (Kostrikov et al., 2021b), a state-of-the-art algorithm for offline RL. We also compare our online finetuning performance to the finetuning variant of IQL, which essentially incorporates the advantage weighted actor critic (AWAC) (Nair et al., 2020) technique for finetuning. For a purely online baseline, we also report the results of soft actor critic (SAC) (Haarnoja et al., 2018a) at 200k online steps. We use the official Pytorch implementation² for DT, the official JAX implementation³ for IQL, and the Pytorch implementation⁴ (Yarats & Kostrikov, 2020) for SAC.

Hyperparameters. We use the default hyperparameters in the open-source codebase for DT and IQL, and those in Haarnoja et al. (2018a) for SAC.⁵ Following the setup of Kostrikov et al. (2021b), the replay buffer we use for IQL and SAC can contain up to 1 million transitions for Gym tasks and 2 million transitions for AntMaze. To make a fair comparison, we restrict the size of the ODT replay buffer so that the maximum number of transitions matches with IQL and SAC, which is 1000 for Gym and 2857 for AntMaze. However, we observe that a smaller replay buffer is beneficial for AntMaze and we use 1500 in our experiments. IQL and SAC both make one gradient update after each online step, and ODT runs 300 gradient updates between two consecutive exploration rollouts. The complete list of hyperparameters of ODT are summarized in Appendix C.

Results and Analysis. For each method, we train 10 instances with different seeds. For each instance, we run 10 evaluation trajectories for Gym tasks and 100 ones for AntMaze tasks, respectively. Table 5.1 report the results.

For simply offline pretraining, IQL outperforms both ODT and DT on most tasks and datasets. However, with an additional budget of 200k online samples, we observe notable performance improvement for ODT on Hopper, Walker, and AntMaze and the relative performance gap δ_{ODT} between the offline and online phase is significant for ODT. In contrast, IQL only shows limited improvements δ_{IQL} given the same online budget. On average, the relative improvements due to finetuning for ODT (δ_{ODT}) are $\sim 9\times$ to those of IQL (δ_{IQL}) for the Gym tasks and datasets. The average absolute performance is similar for both approaches, even though IQL had a much better pretrained initialization to begin with. We also obtain significant improvements due to

²<https://github.com/kzl/decision-transformer>

³https://github.com/ikostrikov/implicit_q_learning

⁴https://github.com/denisyarats/pytorch_sac

⁵The hyperparameters in the Pytorch SAC codebase are optimized for dm-control tasks. We found those used in Haarnoja et al. (2018a) lead to better results.

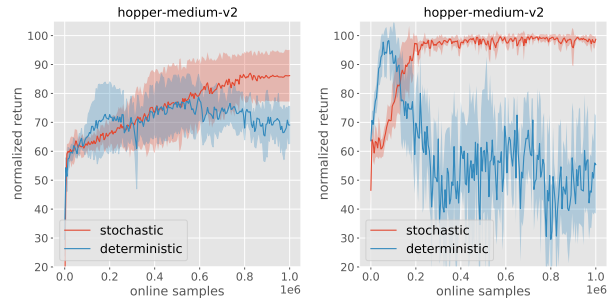


Figure 5.1. Comparison of ODT (red) with a deterministic variant (blue) in terms of training stability on the Hopper benchmark. For both small (L) and large (R) architectures, ODT is stable whereas the performance of the deterministic policy declines and exhibits high variability.

finetuning on the challenging Ant environments and match IQL’s absolute performance. Here, we note that IQL struggles to improve its performance and moreover, its finetuning protocol could also degrade the initial performance (e.g., AntMaze with umaze-diverse dataset).

Finally, it is also instructive to view the results through the lens of online training. We report the performance of SAC, the representative baseline of online RL methods under a sample budget of 200k online interactions. SAC performs substantially worse than ODT on all the Gym tasks except HalfCheetah. In addition, we found that SAC fails to learn non-trivial policies for AntMaze. While it still remains an open question that if transformer-based model-free RL methods can be learned purely online, our results suggest that ODT can significantly benefit in practical regimes with offline data and limited budget for online interactions.

5.2 Ablation Study

We ablate the design choices for ODT to identify the components that are key to its performance. Due to the lack of space, we defer additional experiments to Appendix A.

Stochastic vs Deterministic Policy. Stochasticity is an important component for ODT to enable exploration and stable online training. To demonstrate its effect, we compare ODT to a deterministic variant using the same finetuning framework presented in Section 4. The deterministic variant uses the same architecture as DT, which predicts actions using a fully connected layer at the end and is optimized via the ℓ_2 loss. Figure 5.1 compares the average performance of 5 training instances on two different model architectures. The left panel plots the results on models with small capacity where neither of them solves the environment. The performance of the deterministic variant starts to decrease after collecting 600k online samples. In contrast, ODT is stable and gives higher returns. If we drastically increase the model capacity, as shown in the right panel, both approaches will improve but the performance of deterministic variant quickly degrades and fluctuates, whereas ODT is

Online Decision Transformer

dataset	DT	ODT (offline)	ODT (0.2m)	δ_{ODT}	IQL (offline)	IQL (0.2m)	δ_{IQL}	SAC (0.2m)
hopper-medium	61.03 \pm 5.11	66.95 \pm 3.26	97.54 \pm 2.10	30.59	63.81 \pm 9.15	66.79 \pm 4.07	2.98	66.76 \pm 27.13
hopper-medium-replay	62.75 \pm 15.05	86.64 \pm 5.41	88.89 \pm 6.33	2.25	92.13 \pm 10.43	96.23 \pm 4.35	4.10	
walker2d-medium	72.03 \pm 4.32	72.19 \pm 6.49	76.79 \pm 2.30	4.60	79.89 \pm 3.06	80.33 \pm 2.33	0.44	34.20 \pm 18.43
walker2d-medium-replay	42.53 \pm 15.36	68.92 \pm 4.79	76.86 \pm 4.04	7.94	73.67 \pm 6.37	70.55 \pm 5.81	-3.12	
halfcheetah-medium	42.43 \pm 0.30	42.72 \pm 0.46	42.16 \pm 1.48	-0.56	47.37 \pm 0.29	47.41 \pm 0.15	0.04	55.73 \pm 4.19
halfcheetah-medium-replay	35.92 \pm 1.56	39.99 \pm 0.68	40.42 \pm 1.61	0.43	44.10 \pm 1.14	44.14 \pm 0.3	0.04	
ant-medium	93.56 \pm 4.94	91.33 \pm 4.13	90.79 \pm 5.80	-0.54	99.92 \pm 5.86	100.85 \pm 2.02	0.93	30.03 \pm 7.98
ant-medium-replay	89.08 \pm 5.33	86.56 \pm 3.26	91.57 \pm 2.73	5.01	91.21 \pm 7.27	91.36 \pm 1.47	0.15	
sum			605.02	49.72		597.66	5.56	
antmaze-umaze	53.3 \pm 5.52	53.10 \pm 4.21	88.5 \pm 5.88	35.4	87.1 \pm 2.81	89.5 \pm 5.43	2.4	/
antmaze-umaze-diverse	52.5 \pm 9.89	50.20 \pm 6.69	56.00 \pm 5.69	7.99	64.4 \pm 8.95	56.8 \pm 6.42	-7.6	/
sum			144.5	43.39		146.3	-5.2	

Table 5.1. Comparison of the average normalized returns on Gym and AntMaze tasks. We report the mean and standard deviation over 10 seeds. Blue: The best performance with 200k online samples. Purple: ODT achieves nearly the same performance as IQL.

stable and performs consistently. We also observe that finetuning methods are generally more stable than purely online counterparts, see Appendix E for further details.

RTG Conditioning. As ODT generates return conditioned policies, we examine strategies for specifying the initial RTG token for online exploration and evaluation. We take an offline pretrained ODT model for Hopper using medium dataset. We vary the initial RTG token g_{eval} for evaluation rollouts and compute the return averaged over 100 evaluation trajectories in the left panel of Figure 5.2. Similar to the observation made by Chen et al. (2021), the actual return of the offline pretrained ODT is strongly correlated with g_{eval} , and it saturates at best possible performance even when conditioned on large out-of-distribution RTGs. We built on this observation to study two mechanisms for choosing the initial RTG token g_{online} during online exploration.

Fixed, Scaled Expert RTG. Figure 5.2 suggests that for an offline model, if g_{eval} is large enough, increasing its value further does not change the return significantly. This motivates us to use a large RTG value even for online exploration, going much beyond the maximum achievable returns for the environment. This is applicable to practical situations where we have no knowledge of the expert performance. We examine the returns of ODT with varying values of g_{online} . Figure 5.2 reports the results in its right panel. The returns are stable when g_{online} is set to 1 – 2x of the expert performance but slightly decrease afterwards along with increasing variance. This suggests that one can use high values of g_{online} but not exceptionally large. For all of our experiments, we set g_{online} to twice the expert performance. We note that setting g_{online} as the expert performance results in comparable returns, see Appendix D.

Curriculum RTGs. Another heuristic we have considered stems from curriculum learning (Bengio et al., 2009; 2015; Florensa et al., 2018; Du et al., 2022). Here, we let g_{online} be the non-stationary q -th percentile of the returns of the trajectories stored in the replay buffer. Figure 5.3 reports

the results. The fixed, scaled RTG strategy outperforms this heuristic.

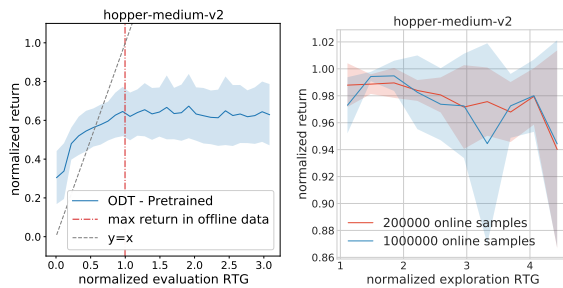


Figure 5.2. The evaluation return obtained by ODT (L) after pre-training when conditioned on various values of g_{eval} , (R) after finetuning with varying values of g_{online} and fixed $g_{\text{eval}} = 3600$.

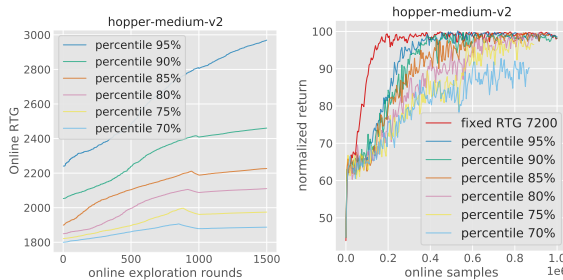


Figure 5.3. Results of an ODT variant that uses the curriculum heuristic to setup the exploration RTG g_{online} . (L) g_{online} for various values of q . (R) Evaluation return comparison. ODT with fixed $g_{\text{online}} = 7200$ (red) outperforms the others, and larger q leads to better performance. Besides, the runs with smaller q collect fewer online samples, this is because the return is highly correlated with the length for dense rewarded environments, see Appendix F.

Hindsight Return Relabeling. Figure 5.4 inspects the return of ODT with and without hindsight return relabeling (Line 6 of Algorithm 2). In the absence of relabeling, ODT

quickly saturates to suboptimal performance and thus, confirming the importance of fixing the RTG tokens prior to appending the trajectory to the training batch of ODT.

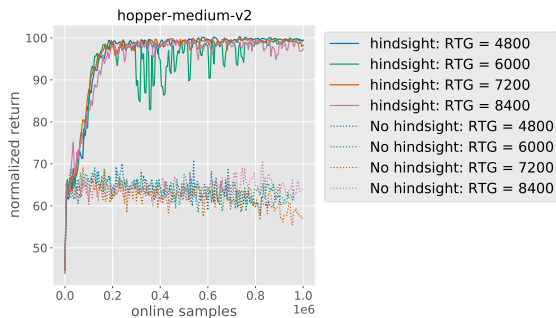


Figure 5.4. Comparison of ODTs trained with hindsight return relabeling enabled or disabled on Hopper using medium dataset. All the other hyperparameters are the same for both methods. Results are averaged over 5 training instances.

6 Discussion

In recent years, significant progress has been made in policy optimization via deep RL algorithms, including policy gradients (Levine & Koltun, 2013; Schulman et al., 2015a; Lillicrap et al., 2015; Schulman et al., 2017), Q-learning (Lange & Riedmiller, 2010; Mnih et al., 2015; Van Hasselt et al., 2016; Gu et al., 2016b; Wang et al., 2016), actor-critic methods (Schulman et al., 2015b; Mnih et al., 2016; Lillicrap et al., 2015; Gu et al., 2016a; Haarnoja et al., 2018a), and model-based learning (Kaiser et al., 2019; Zhang et al., 2019; Lu et al., 2020). Decision Transformer (DT) (Chen et al., 2021), and the closely related work by Janner et al. (2021), provide an alternate perspective by framing offline RL as a sequence modeling problem and solving it via techniques from supervised learning. This provides a simple and scalable framework, including extensions to multi-agent RL (Meng et al., 2021), transfer learning (Boustati et al., 2021), and richer forms of conditioning (Putterman et al.; Furuta et al., 2021).

We proposed ODT, a simple and robust algorithm for fine-tuning a pretrained DT in an online setting, thus further expanding its scope to practical scenarios with a mixture of offline and online interaction data. In the future, it would be interesting to investigate whether a supervised learning approach can account for purely online RL. Our experiment results suggest that ODT is preferred by some but not all the environments. It would be interesting to probe the limits of supervised learning algorithms for RL more broadly, in a similar spirit to Emmons et al. (2021). For example, Kumar et al. (2021) characterize scenarios where RL algorithms outperform classic BC algorithms for offline RL. Similarly, Ortega et al. (2021) point out that sequence modeling approaches for control can create delusions where the agent

mistakes its own actions for task evidence, and propose to treat actions as causal interventions. Finally, we are keen to develop notions of generalization (Kirk et al., 2021; Grover et al., 2018) for ODT and related RL frameworks inspired by supervised learning. We believe these works serve as useful guidance for future work.

Acknowledgement

The authors thank Brandon Amos, Olivier Delalleau, Maryam Fazel-Zarandi, Jiatao Gu, Ilya Kostrikov, Kevin Lu, Mike Rabbat, Shagun Sodhani, Yuandong Tian, Mary Williamson, Lin Xiao and Denis Yarats for insightful discussions.

References

- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., and Zaremba, W. Hindsight Experience Replay. In *Advances in Neural Information Processing Systems*, 2017.
- Bellman, R. A markovian decision process. *Indiana Univ. Math. J.*, 1957.
- Bengio, S., Vinyals, O., Jaitly, N., and Shazeer, N. Scheduled sampling for sequence prediction with recurrent neural networks. *arXiv preprint arXiv:1506.03099*, 2015.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.
- Bottou, L., Peters, J., Quiñero-Candela, J., Charles, D. X., Chikering, D. M., Portugaly, E., Ray, D., Simard, P., and Snelson, E. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(11), 2013.
- Boustati, A., Chockler, H., and McNamee, D. C. Transfer learning with causal counterfactual reasoning in decision transformers. *arXiv preprint arXiv:2110.14355*, 2021.
- Boyd, S., Boyd, S. P., and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners, 2020.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. Decision transformer: Reinforcement learning via sequence modeling. In *Thirty-Fifth Conference on Neural*

- Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=a7APmM4B9d>.
- Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Luan, D., and Sutskever, I. Generative pretraining from pixels. In *International Conference on Machine Learning*, pp. 1691–1703. PMLR, 2020.
- Drusvyatskiy, D. and Xiao, L. Stochastic optimization with decision-dependent distributions. *arXiv preprint arXiv:2011.11173*, 2020.
- Du, Y., Abbeel, P., and Grover, A. It takes four to tango: Multiagent selfplay for automatic curriculum generation. *arXiv preprint arXiv:2202.10608*, 2022.
- Emmons, S., Eysenbach, B., Kostrikov, I., and Levine, S. Rvs: What is essential for offline rl via supervised learning? *arXiv preprint arXiv:2112.10751*, 2021.
- Farahmand, A. M., Munos, R., and Szepesvári, C. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*, 2010.
- Florensa, C., Held, D., Geng, X., and Abbeel, P. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pp. 1515–1528. PMLR, 2018.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- Fujimoto, S. and Gu, S. A minimalist approach to offline reinforcement learning. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=Q32U7dzWXpc>.
- Furuta, H., Matsuo, Y., and Gu, S. S. Generalized decision transformer for offline hindsight information matching. *arXiv preprint arXiv:2111.10364*, 2021.
- Ghosh, D., Gupta, A., Reddy, A., Fu, J., Devin, C., Eysenbach, B., and Levine, S. Learning to reach goals via iterated supervised learning. *arXiv preprint arXiv:1912.06088*, 2019.
- Grover, A., Al-Shedivat, M., Gupta, J. K., Burda, Y., and Edwards, H. Evaluating generalization in multiagent systems using agent-interaction graphs. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1944–1946, 2018.
- Gu, S., Lillicrap, T., Ghahramani, Z., Turner, R. E., and Levine, S. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247*, 2016a.
- Gu, S., Lillicrap, T., Sutskever, I., and Levine, S. Continuous deep q-learning with model-based acceleration. In *International conference on machine learning*, pp. 2829–2838. PMLR, 2016b.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018a.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018b.
- Janner, M., Li, Q., and Levine, S. Offline reinforcement learning as one big sequence modeling problem. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=wgeK563QgSw>.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kirk, R., Zhang, A., Grefenstette, E., and Rocktäschel, T. A survey of generalisation in deep reinforcement learning. *arXiv preprint arXiv:2111.09794*, 2021.
- Kostrikov, I., Fergus, R., Tompson, J., and Nachum, O. Offline reinforcement learning with fisher divergence critic regularization. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5774–5783. PMLR, 18–24 Jul 2021a. URL <https://proceedings.mlr.press/v139/kostrikov21a.html>.
- Kostrikov, I., Nair, A., and Levine, S. Offline reinforcement learning with implicit q-learning, 2021b.
- Kumar, A., Fu, J., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv preprint arXiv:1906.00949*, 2019.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1179–1191. Curran Associates, Inc., 2020a. URL <https://proceedings.neurips.cc/paper/2020/file/>

- 0d2b2061826a5df3221116a5085a6052-Paper.pdf.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020b.
- Kumar, A., Hong, J., Singh, A., and Levine, S. Should i run offline reinforcement learning or behavioral cloning? In *Deep RL Workshop NeurIPS 2021*, 2021.
- Lange, S. and Riedmiller, M. Deep auto-encoder neural networks in reinforcement learning. In *The 2010 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2010.
- Lee, S., Seo, Y., Lee, K., Abbeel, P., and Shin, J. Offline-to-online reinforcement learning via balanced replay and pessimistic q-ensemble. In *5th Annual Conference on Robot Learning*, 2021. URL <https://openreview.net/forum?id=AlJXhEI6J5W>.
- Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Levine, S. and Koltun, V. Guided policy search. In *International conference on machine learning*, pp. 1–9. PMLR, 2013.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Lu, K., Grover, A., Abbeel, P., and Mordatch, I. Reset-free lifelong learning with skill-space planning. *arXiv preprint arXiv:2012.03548*, 2020.
- Lu, K., Grover, A., Abbeel, P., and Mordatch, I. Pretrained transformers as universal computation engines. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022.
- Lu, Y., Hausman, K., Chebotar, Y., Yan, M., Jang, E., Herzog, A., Xiao, T., Irpan, A., Khansari, M., Kalashnikov, D., and Levine, S. AW-opt: Learning robotic skills with imitation and reinforcement at scale. In *5th Annual Conference on Robot Learning*, 2021. URL <https://openreview.net/forum?id=xwEaXgFa0MR>.
- Mendler-Dünnner, C., Perdomo, J. C., Zrnic, T., and Hardt, M. Stochastic optimization for performative prediction. *arXiv preprint arXiv:2006.06887*, 2020.
- Meng, L., Wen, M., Yang, Y., Le, C., Li, X., Zhang, W., Wen, Y., Zhang, H., Wang, J., and Xu, B. Offline pre-trained multi-agent decision transformer: One big sequence model conquers all starcraftii tasks. *arXiv preprint arXiv:2112.02845*, 2021.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937. PMLR, 2016.
- Munos, R. Error bounds for approximate policy iteration. In *ICML*, volume 3, pp. 560–567, 2003.
- Munos, R. Error bounds for approximate value iteration. In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, pp. 1006. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2005.
- Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets. *CoRR*, abs/2006.09359, 2020. URL <https://arxiv.org/abs/2006.09359>.
- Ortega, P. A., Kunesch, M., Deléang, G., Genewein, T., Grau-Moya, J., Veness, J., Buchli, J., Degraeve, J., Piot, B., Perolat, J., et al. Shaking the foundations: delusions in sequence models for interaction and control. *arXiv preprint arXiv:2110.10819*, 2021.
- Perdomo, J., Zrnic, T., Mendler-Dünnner, C., and Hardt, M. Performative prediction. In *International Conference on Machine Learning*, pp. 7599–7609. PMLR, 2020.
- Putterman, A. L., Lu, K., Mordatch, I., and Abbeel, P. Pre-training for language-conditioned imitation with transformers.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.
- Rauber, P., Ummadisingu, A., Mutz, F., and Schmidhuber, J. Hindsight policy gradients. *arXiv preprint arXiv:1711.06006*, 2017.
- Schmidhuber, J. Reinforcement learning upside down: Don’t predict rewards—just map them to actions. *arXiv preprint arXiv:1912.02875*, 2019.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015a.

- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015b.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Srivastava, R. K., Shyam, P., Mutz, F., Jaśkowski, W., and Schmidhuber, J. Training agents using upside-down reinforcement learning. *arXiv preprint arXiv:1912.02877*, 2019.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., and Freitas, N. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, pp. 1995–2003. PMLR, 2016.
- Yarats, D. and Kostrikov, I. Soft actor-critic (sac) implementation in pytorch. https://github.com/denisyarats/pytorch_sac, 2020.
- You, Y., Li, J., Reddi, S., Hseu, J., Kumar, S., Bhojanapalli, S., Song, X., Demmel, J., Keutzer, K., and Hsieh, C.-J. Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962*, 2019.
- Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M., and Levine, S. Solar: Deep structured representations for model-based reinforcement learning. In *International Conference on Machine Learning*, pp. 7444–7453. PMLR, 2019.

A Additional Ablation Study

A.1 Replay Buffer Initialization

We initialize the replay buffer $\mathcal{T}_{\text{replay}}$ by the top trajectories with highest returns in the offline dataset $\mathcal{T}_{\text{offline}}$, see Line 2 in Algorithm 1. Another natural initialization strategy is to randomly select a subset of N trajectories in $\mathcal{T}_{\text{offline}}$. Figure A.1 compares these two strategies in multiple environments. The top N strategy slightly outperforms the random initialization strategy.

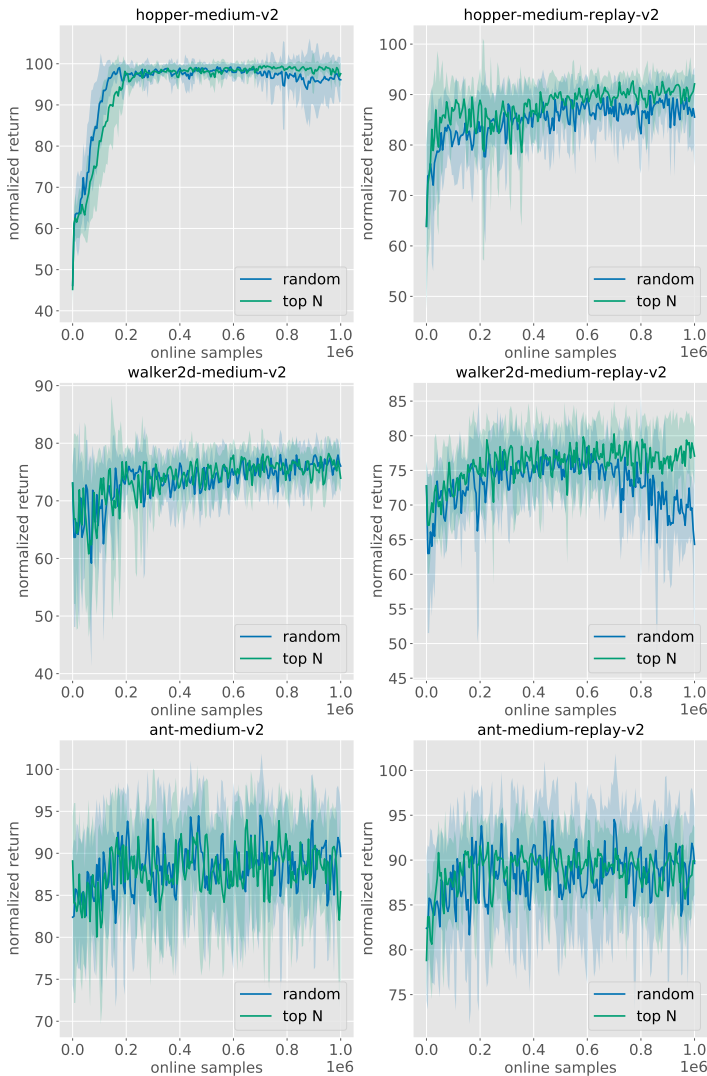


Figure A.1. Comparison of two strategies for initializing the replay buffer $\mathcal{T}_{\text{replay}}$: (1) use the top trajectories with highest returns in the offline dataset $\mathcal{T}_{\text{offline}}$, and (2) randomly select trajectories from $\mathcal{T}_{\text{offline}}$. The size of $\mathcal{T}_{\text{replay}}$ is 1000. For each environment, all the other hyperparameters for both methods are the same as in Section C. Results are averaged over 10 training instances with different seeds.

B Additional Design Choices for ODT

We here discuss the following two configurable components that are critical for the actual performance of ODT. Both of them have an impact on how ODT models the long-horizon dependency for its policy. Experiments demonstrating their influences are also presented below.

Evaluation Context Length As mentioned in Section 3.1, the context length at evaluation time is a hyperparameter we can choose. This parameter adjusts the length of past states and RTGs that the agent’s present action depends on. In the

edge case where the evaluation context length is 1, the agent’s future action only depends on the present state, which means the ODT policies become Markovian, and vice versa. In our experiments, we have found that setting context length 5 at evaluation generally leads to high return across many environments for online finetuning, but the `Ant` task prefers context length 1, see Figure B.1. However, the preferences are mixed for offline training. We summarize the best hyperparameters we have found in Section C.

Positional Embedding In language modeling, positional embeddings are used to equip the input words with their positional information. DT uses the absolute positional embedding which is trained to represent the timesteps of a trajectory, and is added to the embedding of states, actions, and RTGs to help them determine their exact positions within a trajectory. Alternatively, one can use the relative positional embeddings which only account for the order of the states in the input sub-trajectory. However, for problems with non-negative dense rewards, the RTG sequence is monotonically decreasing so that the positional information is partially contained. We have found for some `Gym` tasks in the D4RL benchmark, removing positional embedding improves ODT’s finetuning performance, see Figure B.2 for an example. We conjecture this is due to the fact that the representations of states, actions and RTGs are more disentangled in this case. On the other hand, for goal-reaching tasks with sparse rewards e.g. `AntMaze`, the RTG sequence does not contain timestep information, and positional embedding is needed to stitch the sub-trajectories. Throughout our finetuning experiments in Section 5, we remove the positional embedding for all the `Gym` tasks but use them for the `AntMaze` tasks.

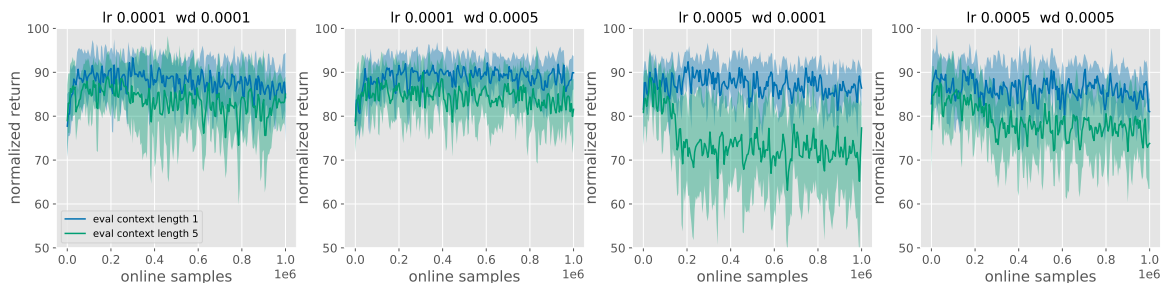


Figure B.1. Comparison of the evaluation context length for finetuning ODT for `Ant` with `medium-replay` offline dataset. For varying values of learning rates and weight decays, context length 1 consistently outperforms 5. Results are averaged over 10 training instances with different seeds.

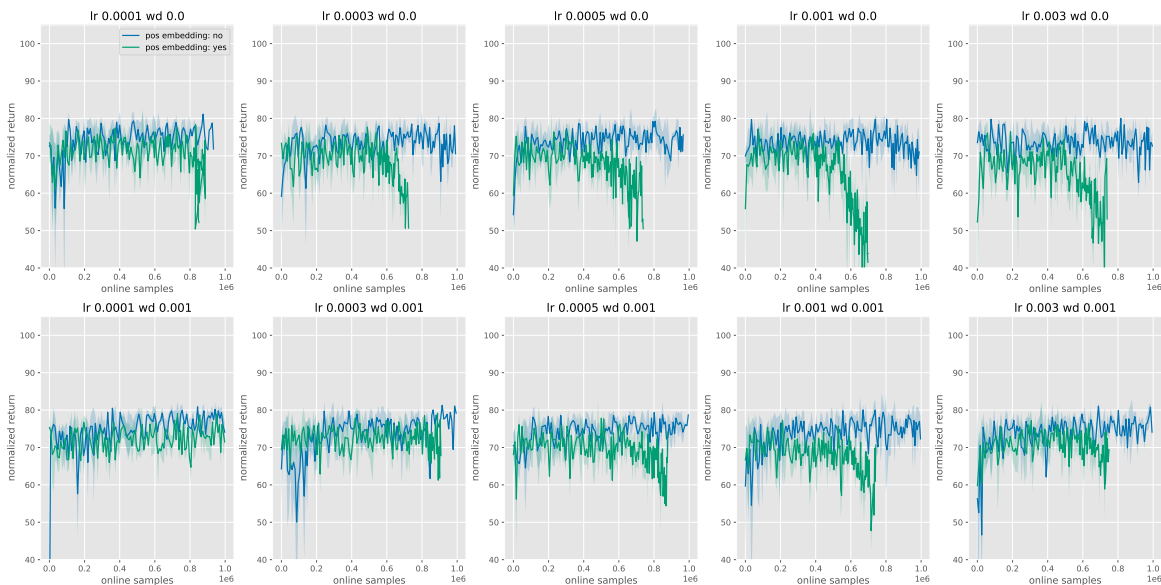


Figure B.2. The online finetuning of ODT for `Walker` with `medium` offline dataset prefers to turn off the positional embedding. For varying values of learning rates and weight decays, training without positional embedding consistently outperforms. Results are averaged over 10 training instances with different seeds.

C Hyperparameters of ODT

We summarize the architecture and other design hyperparameters that we found to work best for ODT in this section. Worthy of note, we use a model with much larger capacity than DT. The transformer we use consists of 4 layers with 4 attention heads, and the embedding dimension is 512. Our intuition is that ODT has more complex structure including the stochasticity to enables online adaption, for which optimization becomes more challenging. Increasing the model capacity might help alleviate the difficulty. Meanwhile, the hyperparameters that result in high return in offline training and online finetuning might not be the same, and we report them separately below.

For all the experiments, we optimize the policy parameter θ by the LAMB optimizer (You et al., 2019), for which we report the learning rates and weight decay parameters below. To optimize the dual parameter λ in problem (7), we optimize the transformed parameter $\log \lambda$ by the Adam optimizer (Kingma & Ba, 2014) with fixed learning rate 0.0001.

As an additional note, the v2 AntMaze datasets we use for our experiments have updated timeout information and we obtain them from the D4RL group. For more results, visit <https://sites.google.com/view/onlinedt/home>.

C.1 Offline Pretraining

hyperparameter	value
Number of layers	4
Number of attention heads	4
Embedding dimension	512
Training Context Length K	20
Dropout	0.1
Nonlinearity function	ReLU
Batch size	256
Learning rate	0.0001
Weight decay	0.001
Gradient norm clip	0.25
Learning rate warmup	linear warmup for 10^4 training steps
Target entropy β	$-\dim(\mathcal{A})$
Total number of updates	10^6

Table C.1. The common hyperparameters that are used to pretrain ODT.

dataset	eval context length	positional embedding	g_{eval}
hopper-medium	20	yes	3600
hopper-medium-replay	20	yes	3600
walker2d-medium	1	no	5000
walker2d-medium-replay	5	no	5000
halfcheetah-medium	5	no	6000
halfcheetah-medium-replay	1	no	6000
ant-medium	5	yes	6000
ant-medium-replay	20	yes	6000
antmaze-umaze	1	yes	1
antmaze-umaze-diverse	1	yes	1

Table C.2. The hyperparameters that we found to work best to pretrain ODT in each domain.

For offline training, all the tasks we consider share most of the hyperparameters except evaluation context length and positional embedding. Following (Haarnoja et al., 2018a;b), we set the target entropy β to be the negative action dimensions. Table C.1 lists the common hyperparameters and Table C.2 lists the domain specific ones. Table C.2 also reports the RTG g_{eval} we use for evaluation rollouts.

C.2 Online Finetuning

Likewise, we report the shared hyperparameters in Table C.3 and domain specific ones in Table C.4. Note that we have found embedding dimension 1024 performs much better than 512 for antmaze-umaze-diverse in finetuning, yet worse in offline training. One of our important finding is that pretraining till convergence might hurt the exploration performance,

and we use much fewer number of pretraining updates than offline models in this scenario.

hyperparameter	value
Number of layers	4
Number of attention heads	4
Embedding dimension	1024 for antmaze-umaze-diverse and 512 for all the other tasks
Training Context Length K	20
Dropout	0.1
Nonlinearity function	ReLU
Batch size	256
Updates between rollouts	300
Gradient norm clip	0.25
Learning rate warmup	linear warmup for 10^4 training steps
Target entropy β	$-\dim(\mathcal{A})$

Table C.3. The common hyperparameters that are used to finetune ODT.

dataset	pretraining updates	buffer size	embedding size	learning rate	weight decay	eval context length	g_{eval}	g_{online}	position
hopper-medium	5000	1000	512	0.0001	0.0005	5	3600	7200	no
hopper-medium-replay	5000	1000	512	0.002	0.0001	5	3600	7200	no
walker2d-medium	10000	1000	512	0.001	0.001	5	5000	10000	no
walker2d-medium-replay	10000	1000	512	0.001	0.001	5	5000	10000	no
halfcheetah-medium	5000	1000	512	0.0001	0.0005	5	6000	12000	no
halfcheetah-medium-replay	5000	1000	512	0.0001	0.0005	5	6000	12000	no
ant-medium	10000	1000	512	0.0015	0.0001	1	6000	12000	no
ant-medium-replay	10000	1000	512	0.0001	0.0005	1	6000	12000	no
antmaze-umaze	7000	1500	512	0.001	0	5	1	2	yes
antmaze-umaze-diverse	7000	1500	1024	0.001	0	5	1	2	yes

Table C.4. The hyperparameters that we use to finetune ODT in each domain.

D Comparison of the Exploration RTG g_{online}

dataset	$g_{online} = 2g^*$	$g_{online} = g^*$
hopper-medium	97.54	95.70
hopper-medium-replay	88.89	86.09
walker2d-medium	76.79	74.26
walker2d-medium-replay	76.86	71.76
halfcheetah-medium	42.16	42.97
halfcheetah-medium-replay	40.42	40.85
ant-medium	90.79	88.94
ant-medium-replay	91.57	90.21
antmaze-umaze	88.5	87.7
antmaze-umaze-diverse	58.19	59.4

Table D.1. The evaluation return obtained by ODT after collection 200k online samples, using different values of exploration RTG g_{online} . The returns are comparable. Results are averaged over 10 training instances with different seeds.

Let g^* denote the expert return. We train ODTs where the initial RTG token for the exploration rollouts are set to g^* and $2g^*$, respectively. Table D.1 reports the return after collecting 200k online samples for all the environments we consider, using the same hyperparameters reported in Section C. The obtained returns are comparable.

E Training Stability Comparison

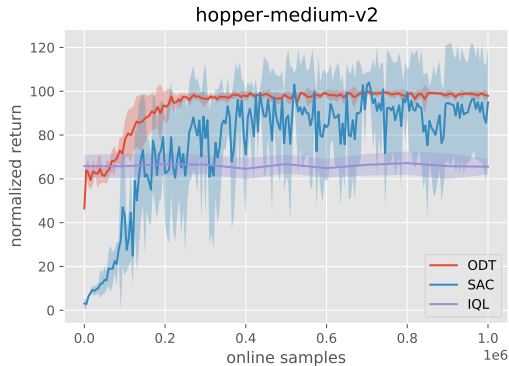


Figure E.1. The training stability comparison of ODT, IQL and SAC. The performance of ODT improves upon pretraining using online interactions, while the performance of IQL stays almost the same as the pretraining result. Both ODT and IQL are stable along the online finetuning yet SAC fluctuates. Results are averaged over 5 training instances with different seeds.

As noted in Section 5.2, we observe that the finetuning methods are generally more stable than purely online counterparts. Figure E.1 plots the results for ODT, IQL and SAC on Hopper with medium dataset. The hyperparameters are set as described in Section 5 for IQL and SAC, and as described in Section C for ODT. We can see that return of ODT improves smoothly over the course of training and the return of IQL keeps stable without significant changes, whereas the return of SAC has high variety although it is increasing.

F Sampling Strategy

The sampling strategy presented in Algorithm 2 has two steps, summarized in Algorithm 3. The first step is to sample a single trajectory τ from the replay buffer $\mathcal{T}_{\text{replay}}$, with probability proportional to its length: $\mathbb{P}_{\mathcal{T}_{\text{replay}}}(\tau) = |\tau| / \sum_{\tau \in \mathcal{T}_{\text{replay}}} |\tau|$. Next, we uniformly sample the subsequences. Figure F.1 plots the sampling probability for each trajectory in the D4RL offline datasets for environments with non-negative dense rewards. We compare it with the importance sampling method, where the probability is proportional to the return of a trajectory. We can see that for many of the datasets, the return of a trajectory is highly correlated with its length.

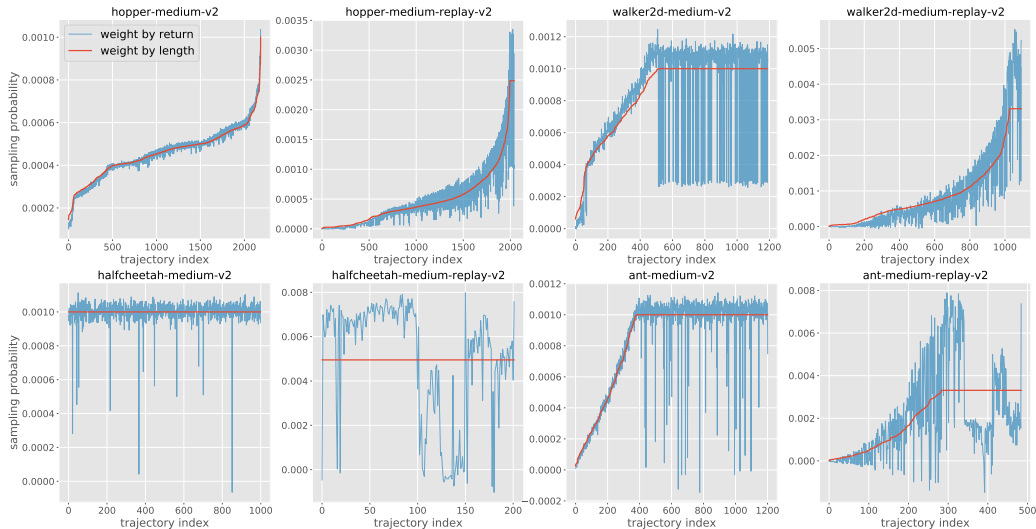


Figure F.1. Comparison of our sampling strategy and importance sampling for the offline datasets collected in environments with non-negative dense rewards.

Algorithm 3: Uniform Subsequence Sampling

- 1 **Input:** replay buffer $\mathcal{T}_{\text{replay}}$, sequence length K .
- 2 Sample a trajectory τ from $\mathcal{T}_{\text{replay}}$ with probability $\mathbb{P}_{\mathcal{T}_{\text{replay}}}(\tau) = |\tau| / \sum_{\tau \in \mathcal{T}_{\text{replay}}} |\tau|$.
- 3 Compute the RTGs $g_t = \sum_{t'=t}^{|\tau|} r_{t'}$, $1 \leq t \leq |\tau|$.
- 4 Sample a timestep $t \in 1, \dots, |\tau|$ uniformly at random.
- 5 $\mathbf{s} \leftarrow (s_t, \dots, s_{\min(t+K-1, |\tau|)})$, $\mathbf{a} \leftarrow (a_t, \dots, a_{\min(t+K-1, |\tau|)})$, $\mathbf{g} \leftarrow (g_t, \dots, g_{\min(t+K-1, |\tau|)})$.
- 6 **Output:** subsequences $(\mathbf{a}, \mathbf{s}, \mathbf{g})$

G Training Dynamics: the convergence of λ

Recall that when optimizing the constrained problem (5), we consider the Lagrangian $L(\theta, \lambda) = J(\theta) + \lambda(\beta - H_{\theta}^T[\mathbf{a}|\mathbf{s}, \mathbf{g}])$ and update θ and λ alternately. As discussed in Section 4.3, under the assumption that the training converges, if λ converges to a value in $[0, 1)$, our overall loss will converge to the desired formulation (9).

In theory, λ can be zero if the constraint is inactive, or positive if the constraint is tight. Figure 4.1 shows an example where λ converges to zero. Here we show another example where λ converges to a positive value between 0 and 1. Figure G.1 shows five example runs for the `walker` environment using `medium` dataset, where we initialize λ by 1, 2, 5, 10, and 100 respectively. The left panel plot the one-sample estimated entropy vs the gradient update iterations, where the first 10000 iterations are offline pretraining, and the later ones are online finetuning. For all the five runs, we can see that the entropy converges to β , which means the constraint $H_{\theta}^T[\mathbf{a}|\mathbf{s}, \mathbf{g}] \geq \beta$ is tight. In this scenario, λ can be positive. The right panel plots the value of λ over the course of training. We can see that λ converges to a same value between 0 and 1 even under different initializations.

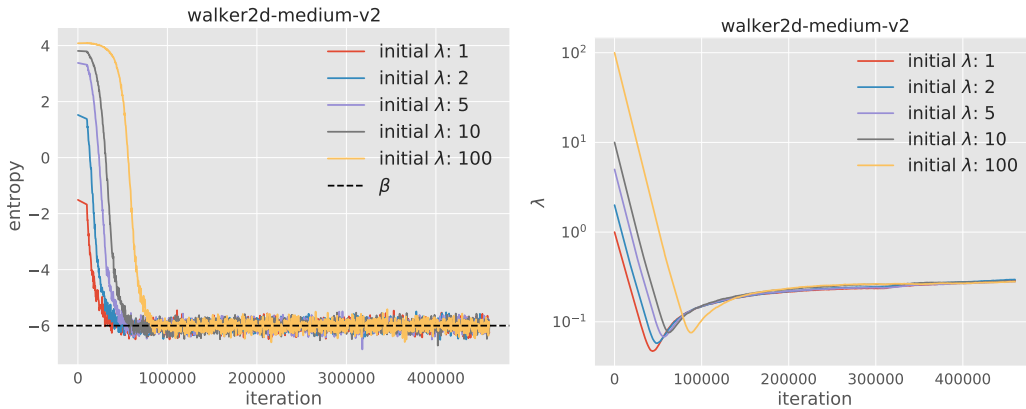


Figure G.1. The estimated entropy and the dual variable λ when optimizing problem (5), including both pretraining and finetuning stages of ODT.

H Negative results

We have tried various heuristics in a few aspects of our algorithms. For modeling, we have tried incorporating both the absolute and relative positional embedding, turning off the dropout regularization, varying the training context length K . For optimization strategies, we have tried the stagewise learning rate decay, the adaptive gradient updates where the number of gradient updates after each exploration rollout is the same as the newly collected trajectory. We found that these heuristics were of limited utility in preliminary experiments and hence we don't include them in the final design.