
Decision Transformer: Reinforcement Learning via Sequence Modeling

Lili Chen^{*1} Kevin Lu^{*1} Aravind Rajeswaran² Kimin Lee¹ Aditya Grover² Michael Laskin¹ Pieter Abbeel¹
Aravind Srinivas^{†1} Igor Mordatch^{†3}

Abstract

We introduce a framework that abstracts Reinforcement Learning (RL) as a sequence modeling problem. This allows us to draw upon the simplicity and scalability of the Transformer architecture, and associated advances in language modeling such as GPT-x and BERT. In particular, we present Decision Transformer, an architecture that casts the problem of RL as conditional sequence modeling. Unlike prior approaches to RL that fit value functions or compute policy gradients, Decision Transformer simply outputs the optimal actions by leveraging a causally masked Transformer. By conditioning an autoregressive model on the desired return (reward), past states, and actions, our Decision Transformer model can generate future actions that achieve the desired return. Despite the simplicity, Decision Transformer matches or exceeds the performance of state-of-the-art model-free offline RL baselines on Atari, OpenAI Gym, and Key-to-Door tasks.

1. Introduction

Recent work has shown transformers (Vaswani et al., 2017) can model high-dimensional distributions of semantic concepts at scale, including effective zero-shot generalization in language (Brown et al., 2020) and out-of-distribution image generation (Ramesh et al., 2021). Given the diversity of successful applications of such models, we seek to examine their application to sequential decision making problems formalized as reinforcement learning (RL). In contrast to prior work using transformers as an architectural choice for components within traditional RL algorithms (Parisotto et al.,

2020; Zambaldi et al., 2018), we seek to study if generative trajectory modeling – i.e. modeling the joint distribution of the sequence of states, actions, and rewards – can serve as a *replacement* for conventional RL algorithms.

We consider the following shift in paradigm: instead of training a policy through conventional RL algorithms like temporal difference (TD) learning (Sutton & Barto, 2018), we will train transformer models on collected experience using a sequence modeling objective. This will allow us to bypass the need for bootstrapping for long term credit assignment – thereby avoiding one of the “deadly triad” (Sutton & Barto, 2018) known to destabilize RL. It also avoids the need for discounting future rewards, as typically done in TD learning, which can induce undesirable short-sighted behaviors. Additionally, we can make use of existing transformer frameworks widely used in language and vision that are easy to scale, utilizing a large body of work studying stable training of transformer models.

In addition to their demonstrated ability to model long sequences, transformers have other advantages. Transformers can perform credit assignment directly via self-attention, in contrast to Bellman backups which slowly propagate rewards and are prone to “distractor” signals (Hung et al., 2019). This can enable transformers to still work effectively in the presence of sparse or distracting rewards. Finally, empirical evidence suggest that a transformer modeling approach can model a wide distribution of behaviors, enabling better generalization and transfer (Ramesh et al., 2021).

We explore our hypothesis by considering *offline RL*, where we will task agents with learning policies from suboptimal data – producing maximally effective behavior from fixed, limited experience. This task is traditionally challenging due to error propagation and value overestimation (Levine et al., 2020). However, it is a natural task when training with a sequence modeling objective. By training an autoregressive model on sequences of states, actions, and returns, we reduce policy sampling to autoregressive generative modeling. We can specify the expertise of the policy – which “skill” to query – by selecting the desired return tokens, acting as a prompt for generation.

^{*}Equal contribution [†]Equal advising ¹UC Berkeley ²Facebook AI Research ³Google Brain. Correspondence to: Lili Chen <lilichen@berkeley.edu>, Kevin Lu <kzl@berkeley.edu>.

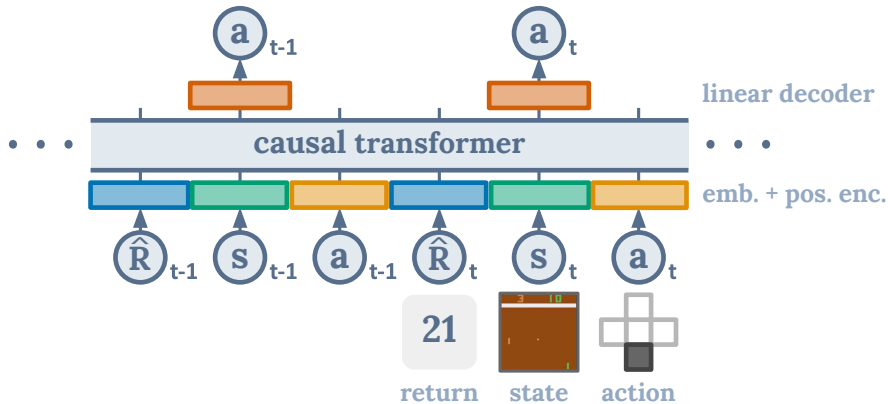


Figure 1. Decision Transformer architecture¹. States, actions, and returns are fed into modality-specific linear embeddings and a positional episodic timestep encoding is added. Tokens are fed into a GPT architecture which predicts actions autoregressively using a causal self-attention mask.




Figure 2. Illustrative example of finding shortest path for a fixed graph (left) posed as reinforcement learning. Training dataset consists of random walk trajectories and their per-node returns-to-go (middle). Conditioned on a starting state and generating largest possible return at each node, Decision Transformer sequences optimal paths.

Illustrative example. To get an intuition for our proposal, consider the task of finding the shortest path on a directed graph, which can be posed as an RL problem. The reward is 0 when the agent is at the goal node and -1 otherwise. We train a GPT (Radford et al., 2018) model to predict next token in a sequence of returns-to-go (sum of future rewards), states, and actions. Training only on random walk data – with no expert demonstrations – we can generate *optimal* trajectories at test time by adding a prior to generate highest possible returns (see more details and empirical results in the Appendix) and subsequently generate the corresponding sequence of actions via conditioning. Thus, by combining the tools of sequence modeling with hindsight return information, we achieve policy improvement without the need for dynamic programming.

Motivated by this observation, we propose Decision Transformer, where we use the GPT architecture to autoregressively model trajectories (shown in Figure 1). We study whether sequence modeling can perform policy optimization by evaluating Decision Transformer on offline RL benchmarks in Atari (Bellemare et al., 2013), OpenAI Gym (Brockman et al., 2016), and Key-to-Door (Mesnard et al., 2020) environments. We show that – *without using dynamic programming* – Decision Transformer matches or ex-

ceeds the performance of state-of-the-art model-free offline RL algorithms (Agarwal et al., 2020; Kumar et al., 2020). Furthermore, in tasks where long-term credit assignment is required, Decision Transformer capably outperforms the RL baselines. With this work, we aim to bridge sequence modeling and transformers with RL, and hope that sequence modeling serves as a strong algorithmic paradigm for RL.

2. Preliminaries

2.1. Offline reinforcement learning

We consider learning in a Markov decision process (MDP) described by the tuple $(\mathcal{S}, \mathcal{A}, P, \mathcal{R})$. The MDP tuple consists of states $s \in \mathcal{S}$, actions $a \in \mathcal{A}$, transition dynamics $P(s'|s, a)$, and a reward function $r = \mathcal{R}(s, a)$. We use s_t , a_t , and $r_t = \mathcal{R}(s_t, a_t)$ to denote the state, action, and reward at timestep t , respectively. A trajectory is made up of a sequence of states, actions, and rewards: $\tau = (s_0, a_0, r_0, s_1, a_1, r_1, \dots, s_T, a_T, r_T)$. The return of a trajectory at timestep t , $R_t = \sum_{t'=t}^T r_{t'}$, is the sum of future rewards from that timestep. The goal in reinforcement learning is to learn a policy which maximizes the expected return $\mathbb{E} \left[\sum_{t=1}^T r_t \right]$ in an MDP. In offline reinforcement learning, instead of obtaining data via environment interactions, we

only have access to some fixed limited dataset consisting of trajectory rollouts of arbitrary policies. This setting is harder as it removes the ability for agents to explore the environment and collect additional feedback.

2.2. Transformers

Transformers were proposed by Vaswani et al. (2017) as an architecture to efficiently model sequential data. These models consist of stacked self-attention layers with residual connections. Each self-attention layer receives n embeddings $\{x_i\}_{i=1}^n$ corresponding to unique input tokens, and outputs n embeddings $\{z_i\}_{i=1}^n$, preserving the input dimensions. The i -th token is mapped via linear transformations to a key k_i , query q_i , and value v_i . The i -th output of the self-attention layer is given by weighting the values v_j by the normalized dot product between the query q_i and other keys k_j :

$$z_i = \sum_{j=1}^n \text{softmax}(\{\langle q_i, k_j \rangle\}_{j=1}^n)_j \cdot v_j. \quad (1)$$

As we shall see later, this allows the layer to assign “credit” by implicitly forming state-return associations via similarity of the query and key vectors (maximizing the dot product). In this work, we use the GPT architecture (Radford et al., 2018), which modifies the transformer architecture with a causal self-attention mask to enable autoregressive generation, replacing the summation/softmax over the n tokens with only the previous tokens in the sequence ($j \in [1, i]$). We defer the other architecture details to the original papers.

3. Method

We present Decision Transformer, which models trajectories autoregressively with minimal modification to the transformer architecture, as summarized in Figure 1.

Trajectory representation. The key desiderata in our choice of trajectory representation are that it should enable transformers to learn meaningful patterns and we should be able to conditionally generate actions at test time. It is nontrivial to model rewards since we would like the model to generate actions based on *future* desired returns, rather than past rewards. As a result, instead of feeding the rewards directly, we feed the model with the returns-to-go $\widehat{R}_t = \sum_{t'=t}^T r_{t'}$. This leads to the following trajectory representation which is amenable to autoregressive training and generation:

$$\tau = \left(\widehat{R}_1, s_1, a_1, \widehat{R}_2, s_2, a_2, \dots, \widehat{R}_T, s_T, a_T \right). \quad (2)$$

At test time, we can specify the desired performance (e.g. 1 for success or 0 for failure), as well as the environment starting state, as the conditioning information to initiate

generation. After executing the generated action for the current state, we decrement the target return by the achieved reward and repeat until episode termination.

Architecture. We feed the last K timesteps into Decision Transformer, for a total of $3K$ tokens (one for each modality: return-to-go, state, or action). To obtain token embeddings, we learn a linear layer for each modality, which projects raw inputs to the embedding dimension, followed by layer normalization (Ba et al., 2016). For environments with visual inputs, the state is fed into a convolutional encoder instead of a linear layer. Additionally, an embedding for each timestep is learned and added to each token – note this is different than the standard positional embedding used by transformers, as one timestep corresponds to three tokens. The tokens are then processed by a GPT (Radford et al., 2018) model, which predicts future action tokens via autoregressive modeling.

Training. We are given a dataset of offline trajectories. We sample minibatches of sequence length K from the dataset. The prediction head corresponding to the input token s_t is trained to predict a_t – either with cross-entropy loss for discrete actions or mean-squared error for continuous actions – and the losses for each timestep are averaged. We did not find predicting the states or returns-to-go to improve performance, although it is easily permissible within our framework (as shown in Section 5.3) and would be an interesting study for future work.

4. Evaluations on Offline RL Benchmarks

In this section, we investigate the performance of Decision Transformer relative to dedicated offline RL and imitation learning algorithms. In particular, our primary points of comparison are model-free offline RL algorithms based on TD-learning, since our Decision Transformer architecture is fundamentally model-free in nature as well. Furthermore, TD-learning is the dominant paradigm in RL for sample efficiency, and also features prominently as a sub-routine in many model-based RL algorithms (Sutton, 1990; Janner et al., 2019). We also compare with behavior cloning and variants, since it also involves a likelihood based policy learning formulation similar to ours. The exact algorithms depend on the environment but our motivations are as follows:

- **TD learning:** most of these methods use an action-space constraint or value pessimism, and will be the most faithful comparison to Decision Transformer, representing standard RL methods. A state-of-the-art model-free method is Conservative Q-Learning (CQL) (Kumar et al., 2020) which serves as our primary comparison. In addition, we also compare against other prior model-free RL algorithms like BEAR (Kumar et al., 2019b) and




Figure 3. Results comparing Decision Transformer (ours) to TD learning (CQL) and behavior cloning across Atari, OpenAI Gym, and Minigrid. On a diverse set of tasks, Decision Transformer performs comparably or better than traditional approaches. Performance is measured by normalized episode return (see text for details).

BRAC (Wu et al., 2019).

- **Imitation learning:** this regime similarly uses supervised losses for training, rather than Bellman backups. We use behavior cloning here, and include a more detailed discussion in Section 5.1.

We evaluate on both discrete (Atari (Bellemare et al., 2013)) and continuous (OpenAI Gym (Brockman et al., 2016)) control tasks. The former involves high-dimensional observation spaces and requires long-term credit assignment, while the latter requires fine-grained continuous control, representing a diverse set of tasks. Our main results are summarized in Figure 3, where we show averaged normalized performance for each domain.

4.1. Atari

The Atari benchmark (Bellemare et al., 2013) is challenging due to its high-dimensional visual inputs and difficulty of credit assignment arising from the delay between actions and resulting rewards. We evaluate our method on 1% of all samples in the DQN-replay dataset as per Agarwal et al. (2020), representing 500 thousand of the 50 million transitions observed by an online DQN agent (Mnih et al., 2015) during training; we report the mean and standard deviation of 3 seeds. We normalize scores based on a professional gamer, following the protocol of Hafner et al. (2020), where 100 represents the professional gamer score and 0 represents a random policy.

We compare to CQL (Kumar et al., 2020), REM (Agarwal et al., 2020), and QR-DQN (Dabney et al., 2018) on four Atari tasks (Breakout, Qbert, Pong, and Seaquest) that are evaluated in Agarwal et al. (2020). We use context lengths of $K = 30$ for Decision Transformer (except $K = 50$ for Pong). We also report the performance of behavior cloning (BC), which utilizes the same network architecture and hyperparameters as Decision Transformer but does not have return-to-go conditioning². For CQL, REM, and QR-

²We also tried using an MLP with $K = 1$ as in prior work, but found this was worse than the transformer.

DQN baselines, we report numbers directly from the CQL and REM papers. We show results in Table 1. Our method is competitive with CQL in 3 out of 4 games and outperforms or matches REM, QR-DQN, and BC on all 4 games.

4.2. OpenAI Gym

In this section, we consider the continuous control tasks from the D4RL benchmark (Fu et al., 2020). We also consider a 2D reacher environment that is not part of the benchmark, and generate the datasets using a similar methodology to the D4RL benchmark. Reacher is a goal-conditioned task and has sparse rewards, so it represents a different setting than the standard locomotion environments (HalfCheetah, Hopper, and Walker). The different dataset settings are described below.

1. Medium: 1 million timesteps generated by a “medium” policy that achieves approximately one-third the score of an expert policy.
2. Medium-Replay: the replay buffer of an agent trained to the performance of a medium policy (approximately 25k-400k timesteps in our environments).
3. Medium-Expert: 1 million timesteps generated by the medium policy concatenated with 1 million timesteps generated by an expert policy.

We compare to CQL (Kumar et al., 2020), BEAR (Kumar et al., 2019b), BRAC (Wu et al., 2019), and AWR (Peng et al., 2019). CQL represents the state-of-the-art in model-free offline RL, an instantiation of TD learning with value pessimism. Score are normalized so that 100 represents an expert policy, as per Fu et al. (2020). CQL numbers are reported from the original paper; BC numbers are run by us; and the other methods are reported from the D4RL paper. Our results are shown in Table 2. Decision Transformer achieves the highest scores in a majority of the tasks and is competitive with the state of the art in the remaining tasks.

³Given that CQL is generally the strongest TD learning method, for Reacher we only run the CQL baseline.

Game	DT (Ours)	CQL	QR-DQN	REM	BC
Breakout	267.5 ± 97.5	211.1	17.1	8.9	138.9 ± 61.7
Qbert	15.4 ± 11.4	104.2	0.0	0.0	17.3 ± 14.7
Pong	106.1 ± 8.1	111.9	18.0	0.5	85.2 ± 20.0
Seaquest	2.5 ± 0.4	1.7	0.4	0.7	2.1 ± 0.3

Table 1. Normalized scores for the 1% DQN-replay Atari dataset. We report the mean and variance across 3 seeds. Best mean scores are highlighted in bold. Decision Transformer performs comparably to CQL on 3 out of 4 games, and usually outperforms other baselines.

Dataset	Environment	DT (Ours)	CQL	BEAR	BRAC-v	AWR	BC
Medium-Expert	HalfCheetah	86.8 ± 1.3	62.4	53.4	41.9	52.7	59.9
Medium-Expert	Hopper	107.6 ± 1.8	111.0	96.3	0.8	27.1	79.6
Medium-Expert	Walker	108.1 ± 0.2	98.7	40.1	81.6	53.8	36.6
Medium-Expert	Reacher	89.1 ± 1.3	30.6	-	-	-	73.3
Medium	HalfCheetah	42.6 ± 0.1	44.4	41.7	46.3	37.4	43.1
Medium	Hopper	67.6 ± 1.0	58.0	52.1	31.1	35.9	63.9
Medium	Walker	74.0 ± 1.4	79.2	59.1	81.1	17.4	77.3
Medium	Reacher	51.2 ± 3.4	26.0	-	-	-	48.9
Medium-Replay	HalfCheetah	36.6 ± 0.8	46.2	38.6	47.7	40.3	4.3
Medium-Replay	Hopper	82.7 ± 7.0	48.6	33.7	0.6	28.4	27.6
Medium-Replay	Walker	66.6 ± 3.0	26.7	19.2	0.9	15.5	36.9
Medium-Replay	Reacher	18.0 ± 2.4	19.0	-	-	-	5.4
Average (Without Reacher)		74.7	63.9	48.2	36.9	34.3	46.4
Average (All Settings)		69.2	54.2	-	-	-	47.7

Table 2. Results for D4RL datasets³. We report the mean and variance for three seeds. Decision Transformer (DT) outperforms conventional RL algorithms on almost all tasks.

5. Discussion

5.1. Does Decision Transformer perform behavior cloning on a subset of the data?

In this section, we seek to gain insight into whether Decision Transformer can be thought of as performing imitation learning on a subset of the data with a certain return. To investigate this, we propose a new method, Percentile Behavior Cloning (%BC), where we run behavior cloning on only the top $X\%$ of timesteps in the dataset, ordered by episode returns. The percentile $X\%$ interpolates between standard BC ($X = 100\%$) that trains on the entire dataset and only cloning the best observed trajectory ($X \rightarrow 0\%$), trading off between better generalization by training on more data with training a specialized model that focuses on a desirable subset of the data.

We show full results comparing %BC to Decision Transformer and CQL in Table 3, sweeping over $X \in [10\%, 25\%, 40\%, 100\%]$. Note that the only way to choose the optimal subset for cloning is to evaluate using rollouts from the environment, so %BC is not a realistic approach; rather, it serves to provide insight into the behavior of Decision Transformer. When data is plentiful – as in the D4RL regime – we find %BC can match or beat other offline RL methods. On most environments, Decision Transformer is competitive with the performance of the best %BC, indicat-

ing it can hone in on a particular subset after training on the entire dataset distribution.

In contrast, when we study low data regimes – such as Atari, where we use 1% of a replay buffer as the dataset – %BC is weak (shown in Table 4). This suggests that in scenarios with relatively low amounts of data, Decision Transformer can outperform %BC by using all trajectories in the dataset to improve generalization, even if those trajectories are dissimilar from the return conditioning target. Our results indicate that Decision Transformer can be more effective than simply performing imitation learning on a subset of the dataset. On the tasks we considered, Decision Transformer either outperforms or is competitive to %BC, without the confound of having to select the optimal subset.

5.2. How well does Decision Transformer model the distribution of returns?

We evaluate the ability of Decision Transformer to understand return-to-go tokens by varying the desired target return over a wide range – evaluating the multi-task distribution modeling capability of transformers. Figure 5 shows the average sampled return accumulated by the agent over the course of the evaluation episode for varying values of target return. On every task, the desired target returns and the true observed returns are highly correlated. On some tasks like

Dataset	Environment	DT (Ours)	10%BC	25%BC	40%BC	100%BC	CQL
Medium	HalfCheetah	42.6 ± 0.1	42.9	43.0	43.1	43.1	44.4
Medium	Hopper	67.6 ± 1.0	65.9	65.2	65.3	63.9	58.0
Medium	Walker	74.0 ± 1.4	78.8	80.9	78.8	77.3	79.2
Medium	Reacher	51.2 ± 3.4	51.0	48.9	58.2	58.4	26.0
Medium-Replay	HalfCheetah	36.6 ± 0.8	40.8	40.9	41.1	4.3	46.2
Medium-Replay	Hopper	82.7 ± 7.0	70.6	58.6	31.0	27.6	48.6
Medium-Replay	Walker	66.6 ± 3.0	70.4	67.8	67.2	36.9	26.7
Medium-Replay	Reacher	18.0 ± 2.4	33.1	16.2	10.7	5.4	19.0
Average		56.1	56.7	52.7	49.4	39.5	43.5

Table 3. Comparison between Decision Transformer (DT) and Percentile Behavior Cloning (%BC).

Game	DT (Ours)	10%BC	25%BC	40%BC	100%BC
Breakout	267.5 ± 97.5	28.5 ± 8.2	73.5 ± 6.4	108.2 ± 67.5	138.9 ± 61.7
Qbert	15.4 ± 11.4	6.6 ± 1.7	16.0 ± 13.8	11.8 ± 5.8	17.3 ± 14.7
Pong	106.1 ± 8.1	2.5 ± 0.2	13.3 ± 2.7	72.7 ± 13.3	85.2 ± 20.0
Seaquest	2.5 ± 0.4	1.1 ± 0.2	1.1 ± 0.2	1.6 ± 0.4	2.1 ± 0.3

Table 4. %BC scores for Atari. We report the mean and variance across 3 seeds. Decision Transformer usually outperforms %BC.

Pong, HalfCheetah and Walker, Decision Transformer generates trajectories that almost perfectly match the desired returns (as indicated by the overlap with the oracle line). Furthermore, on some Atari tasks like Seaquest, we can prompt the Decision Transformer with higher returns than the maximum episode return available in the dataset, demonstrating that Decision Transformer is sometimes capable of extrapolation.




Figure 4. **Left:** Averages of running return probabilities predicted by the transformer model for three types of episode outcomes. **Right:** Transformer attention weights from all timesteps superimposed for a particular successful episode. The model attends to steps near pivotal events in the episode, such as picking up the key and reaching the door.

5.3. Does Decision Transformer perform effective long-term credit assignment?

To evaluate long-term credit assignment capabilities of our model, we consider a variant of the Key-to-Door environment proposed in Mesnard et al. (2020). This is a grid-based environment with a sequence of three phases: (1) in the first phase, the agent is placed in a room with a key; (2) then, the agent is placed in an empty room; (3) and finally, the agent is placed in a room with a door. The agent receives

a binary reward when reaching the door in the third phase, but **only** if it picked up the key in the first phase. We train on datasets of trajectories generated by applying random actions and report success rates in Table 5. Methods that use hindsight return information: our Decision Transformer model and %BC (trained only on successful episodes) are able to learn effective policies – producing near-optimal paths, despite only training on random walks. TD learning (CQL) cannot effectively propagate Q-values over the long horizons involved and gets poor performance.

5.4. Can transformers be accurate critics in sparse reward settings?

In previous sections, we established that decision transformer can produce effective policies (actors). We now evaluate whether transformer models can also be effective critics. We modify Decision Transformer to output return tokens in addition to action tokens on the Key-to-Door environment. We find that the transformer continuously updates reward probability based on events during the episode, shown in Figure 4 (Left). Furthermore, we find the transformer attends to critical events in the episode (picking up the key or reaching the door), shown in Figure 4 (Right), indicating formation of state-reward associations as discussed in Raposo et al. (2021) and enabling accurate value prediction.

5.5. Does Decision Transformer perform well in sparse reward settings?

A known weakness of TD learning algorithms is that they require densely populated rewards, which can be unrealistic and/or expensive. In contrast, Decision Transformer can improve robustness in these settings since it makes minimal

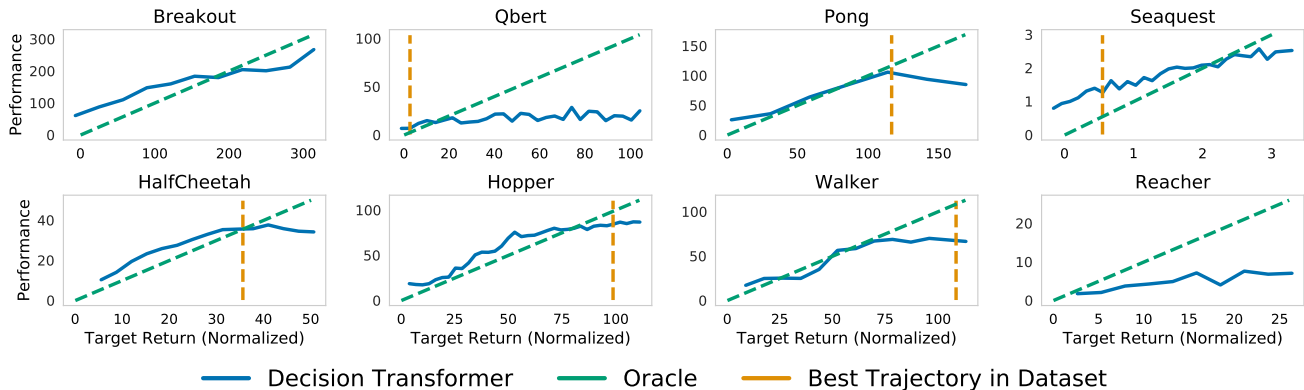


Figure 5. Sampled (evaluation) returns accumulated by Decision Transformer when conditioned on the specified target (desired) returns. **Top:** Atari. **Bottom:** D4RL medium-replay datasets.

Dataset	DT (Ours)	CQL	BC	%BC	Random
1K Random Trajectories	71.8%	13.1%	1.4%	69.9%	3.1%
10K Random Trajectories	94.6%	13.3%	1.6%	95.1%	3.1%

Table 5. Success rate for Key-to-Door environment. Methods using hindsight (Decision Transformer, %BC) can learn successful policies, while TD learning struggles to perform credit assignment.

assumptions on the density of the reward. To evaluate this, we consider a delayed return version of the D4RL benchmarks where the agent does not receive any rewards along the trajectory, and instead receives the cumulative reward of the trajectory in the final timestep. Our results for delayed returns are shown in Table 6. Delayed returns minimally affect Decision Transformer; and due to the nature of the training process, while imitation learning methods are reward agnostic. While TD learning collapses, Decision Transformer and %BC perform well, indicating that Decision Transformer can be more robust to delayed rewards.

See Appendix D for additional discussions.

6. Related Work

Offline reinforcement learning. To mitigate the impact of distribution shift in offline RL, prior algorithms either (a) constrain the policy action space (Fujimoto et al., 2019; Kumar et al., 2019a; Siegel et al., 2020) or (b) incorporate value pessimism (Fujimoto et al., 2019; Kumar et al., 2020), or (c) incorporate pessimism into learned dynamics models (Kidambi et al., 2020; Yu et al., 2020). Since we do not use Decision Transformers to explicitly learn the dynamics model, we primarily compare against model-free algorithms in our work; in particular, adding a dynamics model tends to improve the performance of model-free algorithms. Another line of work explores learning wide behavior distribution from an offline dataset by learning a task-agnostic set of skills, either with likelihood-based approaches (Ajay et al., 2020; Campos et al., 2020; Pertsch et al., 2020; Singh et al.,

2021) or by maximizing mutual information (Eysenbach et al., 2019; Lu et al., 2020; Sharma et al., 2020). Our work is similar to the likelihood-based approaches, which do not use iterative Bellman updates – although we use a simpler sequence modeling objective instead of a variational method, and use rewards for conditional generation of behaviors.

Supervised learning in reinforcement learning settings.

Some prior methods for reinforcement learning bear more resemblance to static supervised learning, such as Q-learning (Watkins, 1989; Mnih et al., 2013), which still uses iterative backups, or likelihood-based methods such as behavior cloning, which do not (discussed in previous section). Recent work (Srivastava et al., 2019; Kumar et al., 2019c; ogm, 2019) studies “upside-down” reinforcement learning (UDRL), which are similar to our method in seeking to model behaviors with a supervised loss conditioned on the target return. A key difference in our work is the shift of motivation to sequence modeling rather than supervised learning: while the practical methods differ primarily in the context length and architecture, sequence modeling enables behavior modeling even without access to the reward, in a similar style to language (Radford et al., 2018) or images (Chen et al., 2020), and is known to scale well (Brown et al., 2020). The method proposed by Kumar et al. (2019c) is most similar to our method with $K = 1$, which we find sequence modeling/long contexts to outperform (see Section D). Ghosh et al. (2019) extends prior UDRL methods to use state goal conditioning, rather than rewards, and Paster et al. (2020) further use an LSTM with state goal conditioning for goal-conditioned online RL settings. Concurrent to our

Dataset	Environment	Delayed (Sparse)		Agnostic		Original (Dense)	
		DT (Ours)	CQL	BC	%BC	DT (Ours)	CQL
Medium-Expert	Hopper	107.3 ± 3.5	9.0	59.9	102.6	107.6	111.0
Medium	Hopper	60.7 ± 4.5	5.2	63.9	65.9	67.6	58.0
Medium-Replay	Hopper	78.5 ± 3.7	2.0	27.6	70.6	82.7	48.6

Table 6. Results for D4RL datasets with delayed (sparse) reward. Decision Transformer (DT) and imitation learning are minimally affected by the removal of dense rewards, while CQL fails.

work, Janner et al. (2021) propose Trajectory Transformer, which is similar to Decision Transformer but additionally uses state and return prediction, as well as discretization, which incorporates model-based components. We believe that their experiments, in addition to our results, highlight the potential for sequence modeling to be a generally applicable idea for reinforcement learning.

Credit assignment. Many works have studied better credit assignment via state-association, learning an architecture which decomposes the reward function such that certain “important” states comprise most of the credit (Ferret et al., 2019; Harutyunyan et al., 2019; Mesnard et al., 2020). They use the learned reward function to change the reward of an actor-critic algorithm to help propagate signal over long horizons. Similar to our long-term setting, some works have specifically shown such state-associative architectures can perform better in delayed reward settings (Arjona-Medina et al., 2018; Hung et al., 2019; Liu et al., 2019; Raposo et al., 2021). In contrast, we allow these properties to naturally emerge in a transformer architecture, without having to explicitly learn a reward function or a critic.

Conditional language generation. Various works have studied guided generation for images (Karras et al., 2019) and language (Ghazvininejad et al., 2017; Weng, 2021). Several works (Ficler & Goldberg, 2017; Hu et al., 2017; Rajani et al., 2019; Yu et al., 2017; Ziegler et al., 2019; Keskar et al., 2019) have explored training or fine-tuning of models for controllable text generation. Class-conditional language models can also be used to learn discriminators to guide generation (Dathathri et al., 2019; Ghazvininejad et al., 2017; Holtzman et al., 2018; Krause et al., 2020). However, these approaches mostly assume constant “classes”, while in reinforcement learning the reward signal is time-varying. Furthermore, it is more natural to prompt the model desired target return and continuously decrease it by the observed rewards over time, since the transformer model and environment jointly generate the trajectory.

Attention and transformer models. Transformers (Vaswani et al., 2017) have been applied successfully to many tasks in natural language processing (Devlin et al., 2018; Radford et al., 2018) and computer vision (Carion et al., 2020; Dosovitskiy et al., 2020). However, transformers are relatively unstudied in RL, mostly due to differing

nature of the problem, such as higher variance in training. Zambaldi et al. (2018) showed that augmenting transformers with relational reasoning improve performance in combinatorial environments and Ritter et al. (2020) showed iterative self-attention allowed for RL agents to better utilize episodic memories. Parisotto et al. (2020) discussed design decisions for more stable training of transformers in the high-variance RL setting. Unlike our work, these still use actor-critic algorithms for optimization, focusing on novelty in architecture. Additionally, in imitation learning, some works have studied transformers as a replacement for LSTMs: Dasari & Gupta (2020) study one-shot imitation learning, and Abramson et al. (2020) combine language and image modalities for text-conditioned behavior generation.

7. Conclusion

We proposed Decision Transformer, seeking to unify ideas in language/sequence modeling and reinforcement learning. On standard offline RL benchmarks, we showed Decision Transformer can match or outperform strong algorithms designed explicitly for offline RL with minimal modifications from standard language modeling architectures.

We hope this work inspires investigation into using large transformer models for RL. We used a simple supervised loss that was effective in our experiments, but applications to large-scale datasets could benefit from self-supervised pretraining tasks. In addition, one could consider more sophisticated embeddings for returns, states, and actions – for instance, conditioning on return distributions to model stochastic settings instead of deterministic returns. Transformer models can also be used to model the state evolution of trajectory, potentially serving as an alternative to model-based RL, and we hope to explore this in future work.

Successful real-world deployment of AI systems requires strong generalization to unseen situations. While we do not directly evaluate generalization ability, we utilize a method known to generalize well in domains such as language and vision. Given our promising results on offline RL benchmarks, we are excited about the potential of larger RL systems built upon the Decision Transformer framework. For real-world applications, it will also be important to analyze the errors transformers make in MDP settings and biases in the datasets used for training.

References

- Acting without rewards. 2019. URL <https://ogma.ai/2019/08/acting-without-rewards/>.
- Abramson, J., Ahuja, A., Barr, I., Brussee, A., Carnevale, F., Cassin, M., Chhapparia, R., Clark, S., Damoc, B., Dudzik, A., et al. Imitating interactive intelligence. *arXiv preprint arXiv:2012.05672*, 2020.
- Agarwal, R., Schuurmans, D., and Norouzi, M. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, 2020.
- Ajay, A., Kumar, A., Agrawal, P., Levine, S., and Nachum, O. Opal: Offline primitive discovery for accelerating offline reinforcement learning. *arXiv preprint arXiv:2010.13611*, 2020.
- Arjona-Medina, J. A., Gillhofer, M., Widrich, M., Unterthiner, T., Brandstetter, J., and Hochreiter, S. Ruder: Return decomposition for delayed rewards. *arXiv preprint arXiv:1806.07857*, 2018.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Campos, V., Trott, A., Xiong, C., Socher, R., Giro-i Nieto, X., and Torres, J. Explore, discover and learn: Unsupervised discovery of state-covering skills. In *International Conference on Machine Learning*, 2020.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 2020.
- 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.
- Dabney, W., Rowland, M., Bellemare, M., and Munos, R. Distributional reinforcement learning with quantile regression. In *Conference on Artificial Intelligence*, 2018.
- Dasari, S. and Gupta, A. Transformers for one-shot visual imitation. *arXiv preprint arXiv:2011.05970*, 2020.
- Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J., and Liu, R. Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., and Clune, J. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations*, 2019.
- Ferret, J., Marinier, R., Geist, M., and Pietquin, O. Self-attentional credit assignment for transfer in reinforcement learning. *arXiv preprint arXiv:1907.08027*, 2019.
- Ficler, J. and Goldberg, Y. Controlling linguistic style aspects in neural language generation. *arXiv preprint arXiv:1707.02633*, 2017.
- 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., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, 2019.
- Gao, Y., Xu, H., Lin, J., Yu, F., Levine, S., and Darrell, T. Reinforcement learning from imperfect demonstrations. *arXiv preprint arXiv:1802.05313*, 2018.
- Ghazvininejad, M., Shi, X., Priyadarshi, J., and Knight, K. Hafez: an interactive poetry generation system. In *Proceedings of ACL, System Demonstrations*, 2017.
- Ghosh, D., Gupta, A., Fu, J., Reddy, A., Devin, C., Eysenbach, B., and Levine, S. Learning to reach goals without reinforcement learning. *arXiv preprint arXiv:1912.06088*, 2019.

- Hafner, D., Lillicrap, T., Norouzi, M., and Ba, J. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- Harutyunyan, A., Dabney, W., Mesnard, T., Azar, M., Piot, B., Heess, N., van Hasselt, H., Wayne, G., Singh, S., Precup, D., et al. Hindsight credit assignment. *arXiv preprint arXiv:1912.02503*, 2019.
- Holtzman, A., Buys, J., Forbes, M., Bosselut, A., Golub, D., and Choi, Y. Learning to write with cooperative discriminators. *arXiv preprint arXiv:1805.06087*, 2018.
- Hu, Z., Yang, Z., Liang, X., Salakhutdinov, R., and Xing, E. P. Toward controlled generation of text. In *International Conference on Machine Learning*, 2017.
- Hung, C.-C., Lillicrap, T., Abramson, J., Wu, Y., Mirza, M., Carnevale, F., Ahuja, A., and Wayne, G. Optimizing agent behavior over long time scales by transporting value. *Nature communications*, 10(1):1–12, 2019.
- Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, pp. 12498–12509, 2019.
- Janner, M., Li, Q., and Levine, S. Reinforcement learning as one big sequence modeling problem. *arXiv preprint arXiv:2106.02039*, 2021.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition*, 2019.
- Keskar, N. S., McCann, B., Varshney, L. R., Xiong, C., and Socher, R. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*, 2019.
- Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. Morel: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- Krause, B., Gotmare, A. D., McCann, B., Keskar, N. S., Joty, S., Socher, R., and Rajani, N. F. Gedi: Generative discriminator guided sequence generation. *arXiv preprint arXiv:2009.06367*, 2020.
- Kumar, A., Fu, J., Soh, M., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, 2019a.
- Kumar, A., Fu, J., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv preprint arXiv:1906.00949*, 2019b.
- Kumar, A., Peng, X. B., and Levine, S. Reward-conditioned policies. *arXiv preprint arXiv:1912.13465*, 2019c.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Liu, Y., Luo, Y., Zhong, Y., Chen, X., Liu, Q., and Peng, J. Sequence modeling of temporal credit assignment for episodic reinforcement learning. *arXiv preprint arXiv:1905.13420*, 2019.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Lu, K., Grover, A., Abbeel, P., and Mordatch, I. Reset-free lifelong learning with skill-space planning. *arXiv preprint arXiv:2012.03548*, 2020.
- Mesnard, T., Weber, T., Viola, F., Thakoor, S., Saade, A., Harutyunyan, A., Dabney, W., Stepleton, T., Heess, N., Guez, A., et al. Counterfactual credit assignment in model-free reinforcement learning. *arXiv preprint arXiv:2011.09464*, 2020.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- 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): 529–533, 2015.
- Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- Parisotto, E., Song, F., Rae, J., Pascanu, R., Gulcehre, C., Jayakumar, S., Jaderberg, M., Kaufman, R. L., Clark, A., Noury, S., et al. Stabilizing transformers for reinforcement learning. In *International Conference on Machine Learning*, 2020.
- Paster, K., McIlraith, S. A., and Ba, J. Planning from pixels using inverse dynamics models. *arXiv preprint arXiv:2012.02419*, 2020.
- Peng, X. B., Kumar, A., Zhang, G., and Levine, S. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

- Pertsch, K., Lee, Y., and Lim, J. J. Accelerating reinforcement learning with learned skill priors. *arXiv preprint arXiv:2010.11944*, 2020.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. 2018.
- Rajani, N. F., McCann, B., Xiong, C., and Socher, R. Explain yourself! leveraging language models for common-sense reasoning. *arXiv preprint arXiv:1906.02361*, 2019.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., and Sutskever, I. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021.
- Raposo, D., Ritter, S., Santoro, A., Wayne, G., Weber, T., Botvinick, M., van Hasselt, H., and Song, F. Synthetic returns for long-term credit assignment. *arXiv preprint arXiv:2102.12425*, 2021.
- Ritter, S., Faulkner, R., Sartran, L., Santoro, A., Botvinick, M., and Raposo, D. Rapid task-solving in novel environments. *arXiv preprint arXiv:2006.03662*, 2020.
- Sharma, A., Gu, S., Levine, S., Kumar, V., and Hausman, K. Dynamics-aware unsupervised discovery of skills. In *International Conference on Learning Representations*, 2020.
- Siegel, N. Y., Springenberg, J. T., Berkenkamp, F., Abdolmaleki, A., Neunert, M., Lampe, T., Hafner, R., and Riedmiller, M. Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. In *International Conference on Learning Representations*, 2020.
- Singh, A., Liu, H., Zhou, G., Yu, A., Rhinehart, N., and Levine, S. Parrot: Data-driven behavioral priors for reinforcement learning. In *International Conference on Learning Representations*, 2021.
- 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.
- Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *ICML*, 1990.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT Press, 2018.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- Watkins, C. Learning from delayed rewards. 01 1989.
- Weng, L. Controllable neural text generation. [lilianweng.github.io/lil-log, 2021. URL https://lilianweng.github.io/lil-log/2021/01/02/controllable-neural-text-generation.html.](https://lilianweng.github.io/lil-log/2021/01/02/controllable-neural-text-generation.html)
- Wolf, T., Chaumond, J., Debut, L., Sanh, V., Delangue, C., Moi, A., Cistac, P., Funtowicz, M., Davison, J., Shleifer, S., et al. Transformers: State-of-the-art natural language processing. In *Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020.
- Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI conference on artificial intelligence*, 2017.
- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C., and Ma, T. Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, 2020.
- Zambaldi, V., Raposo, D., Santoro, A., Bapst, V., Li, Y., Babuschkin, I., Tuyls, K., Reichert, D., Lillicrap, T., Lockhart, E., et al. Deep reinforcement learning with relational inductive biases. In *International Conference on Learning Representations*, 2018.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., and Irving, G. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

A. Experimental Details

A.1. Atari

We build our Decision Transformer implementation for Atari games off of minGPT (<https://github.com/karpathy/minGPT>), a publicly available re-implementation of GPT. We use most of the default hyperparameters from their character-level GPT example (https://github.com/karpathy/minGPT/blob/master/play_char.ipynb). We reduce the batch size (except in Pong), block size, number of layers, attention heads, and embedding dimension for faster training. For processing the observations, we use the DQN encoder from Mnih et al. (2015) with an additional linear layer to project to the embedding dimension.

For return-to-go conditioning, we use either $1\times$ or $5\times$ the maximum return in the dataset, but more possibilities exist for principled return-to-go conditioning. In Atari experiments, we use Tanh instead of LayerNorm (as described in Section 3) after embedding each modality, but did this does not make a significant difference in performance. The full list of hyperparameters can be found in Table 7.

A.2. OpenAI Gym

A.2.1. DECISION TRANSFORMER

Our code is based on the Huggingface Transformers library (Wolf et al., 2020). Our hyperparameters on all OpenAI Gym tasks are shown below in Table 8. Heuristically, we find using larger models helps to model the distribution of returns, compared to standard RL model sizes (which learn one policy). For Reacher we use a smaller context length than the other environments, which we find to be helpful as the environment is goal-conditioned and the episodes are shorter. We choose return targets based on expert performance for each environment, except for HalfCheetah where we find 50% performance to be better due to the datasets containing lower relative returns to the other environments. Models were trained for 10^5 gradient steps using the AdamW optimizer (Loshchilov & Hutter, 2017) following PyTorch defaults.

A.2.2. BEHAVIOR CLONING

As briefly mentioned in Section 4.2, we found previously reported behavior cloning baselines to be weak, and so run them ourselves using a similar setup as Decision Transformer. We tried using a transformer architecture, but found using an MLP (as in previous work) to be stronger. We train for 2.5×10^4 gradient steps; training more did not improve performance. Other hyperparameters are shown in Table 9. The percentile behavior cloning experiments use the same hyperparameters.

B. Graph Shortest Path

We give details of the illustrative example discussed in the introduction. The task is to find the shortest path on a fixed directed graph, which can be formulated as an MDP where reward is 0 when the agent is at the goal node and -1 otherwise. The observation is the integer index of the graph node the agent is in. The action is the integer index of the graph node to move to next. The transition dynamics transport the agent to the action’s node index if there is an edge in the graph, while the agent remains at the past node otherwise. The returns-to-go in this problem correspond to negative path lengths and maximizing them corresponds to generating shortest paths.

In this environment, we use the GPT model as described in Section 3 to generate both actions and return-to-go tokens. This makes it possible for the model to generate its own (realizable) returns-to-go \hat{R} . Since we require a return prompt to generate actions and we do assume knowledge of the optimal path length upfront, we use a simple prior over returns that favors shorter paths: $P_{\text{prior}}(\hat{R} = k) \propto T + 1 - k$, where T is the maximum trajectory length. Then, it is combined with the return probabilities generated by the GPT model: $P(\hat{R}_t | s_{0:t}, a_{0:t-1}, \hat{R}_{0:t-1}) = P_{\text{GPT}}(\hat{R}_t | s_{0:t}, a_{0:t-1}, \hat{R}_{0:t-1}) \times P_{\text{prior}}(\hat{R}_t)^{10}$. Note that the prior and return-to-go predictions are entirely computable by the model, and thus avoids the need for any external or oracle information like the optimal path length. Adjustment of generation by a prior has also been used for similar purposes in controllable text generation in prior work (Krause et al., 2020).

We train on a dataset of 1,000 graph random walk trajectories of $T = 10$ steps each with a random graph of 20 nodes and edge sparsity coefficient of 0.1. We report the results in Figure 6, where we find that transformer model is able to significantly improve upon the number of steps required to reach the goal, closely matching performance of optimal paths.

There are two reasons for the favorable performance on this task. In one case, the training dataset of random walk trajectories may contain a segment that directly corresponds to the desired shortest path, in which case it will be generated by the model. In the second case, generated paths are entirely original and are not subsets of trajectories in the training dataset - they are generated from stitching sub-optimal segments. We find this case accounts for 15.8% of generated paths in the experiment.

While this is a simple example and uses a prior on generation that we do not use in other experiments for simplicity, it illustrates how hindsight return information can be used with generation priors to avoid the need for explicit dynamic programming.

Table 7. Hyperparameters of DT (and %BC) for Atari experiments.

Hyperparameter	Value
Number of layers	6
Number of attention heads	8
Embedding dimension	128
Batch size	512 Pong 128 Breakout, Qbert, Seaquest
Context length K	50 Pong 30 Breakout, Qbert, Seaquest
Return-to-go conditioning	90 Breakout ($\approx 1 \times$ max in dataset) 2500 Qbert ($\approx 5 \times$ max in dataset) 20 Pong ($\approx 1 \times$ max in dataset) 1450 Seaquest ($\approx 5 \times$ max in dataset) GeLU, otherwise
Encoder channels	32, 64, 64
Encoder filter sizes	$8 \times 8, 4 \times 4, 3 \times 3$
Encoder strides	4, 2, 1
Max epochs	5
Dropout	0.1
Learning rate	$6 * 10^{-4}$
Adam betas	(0.9, 0.95)
Grad norm clip	1.0
Weight decay	0.1
Learning rate decay	Linear warmup and cosine decay (see code for details)
Warmup tokens	$512 * 20$
Final tokens	$2 * 500000 * K$

C. Atari Task Scores

Table 10 shows the normalized scores used for normalization used in Hafner et al. (2020). Tables 11 and 12 show the raw scores corresponding to Tables 1 and 4, respectively. For %BC scores, we use the same hyperparameters as Decision Transformer for fair comparison. For REM and QR-DQN, there is a slight discrepancy between Agarwal et al. (2020) and Kumar et al. (2020); we report raw data provided to us by REM authors.

D. Additional Discussions

D.1. Context Length Ablation

To assess the importance of access to previous states, actions, and returns, we ablate on the context length K . This is interesting since it is generally considered that the previous state (i.e. $K = 1$) is enough for reinforcement learning algorithms when frame stacking is used, as we do. Table 13 shows that performance of Decision Transformer is significantly worse when $K = 1$, indicating that past information is useful for Atari games. One hypothesis is that when we are representing a distribution of policies – like with sequence modeling – the context allows the transformer to identify *which* policy generated the actions, enabling better learning and/or improving the training dynamics.

D.2. Why does Decision Transformer avoid the need for value pessimism or behavior regularization?

One key difference between Decision Transformer and prior offline RL algorithms is that we do not require policy regularization or conservatism to achieve good performance. Our conjecture is that TD-learning based algorithms learn an approximate value function and improve the policy by optimizing this value function. This act of optimizing a learned function can exacerbate and exploit any inaccuracies in the value function approximation, causing failures in policy improvement. Since Decision Transformer does not require explicit optimization using learned functions as objectives, it avoids the need for regularization or conservatism.

D.3. How can Decision Transformer benefit online RL?

Offline RL and the ability to model behaviors has the potential to enable sample-efficient online RL for downstream tasks. Works studying the transition from offline to online generally find that likelihood-based approaches, like our sequence modeling objective, are more successful (Gao et al., 2018; Nair et al., 2020). As a result, although we studied offline RL in this work, we believe Decision Transformer can meaningfully improve online RL methods by serving as a strong model for behavior generation. For instance, Decision Transformer can serve as a powerful “memorization engine” and in conjunction with powerful exploration

Decision Transformer: Reinforcement Learning via Sequence Modeling

Table 8. Hyperparameters of Decision Transformer for OpenAI Gym experiments.

Hyperparameter	Value
Number of layers	3
Number of attention heads	1
Embedding dimension	128
Nonlinearity function	ReLU
Batch size	64
Context length K	20 HalfCheetah, Hopper, Walker 5 Reacher
Return-to-go conditioning	6000 HalfCheetah 3600 Hopper 5000 Walker 50 Reacher
Dropout	0.1
Learning rate	10^{-4}
Grad norm clip	0.25
Weight decay	10^{-4}
Learning rate decay	Linear warmup for first 10^5 training steps

Table 9. Hyperparameters of Behavior Cloning for OpenAI Gym experiments.

Hyperparameter	Value
Number of layers	3
Embedding dimension	256
Nonlinearity function	ReLU
Batch size	64
Dropout	0.1
Learning rate	10^{-4}
Weight decay	10^{-4}
Learning rate decay	Linear warmup for first 10^5 training steps

algorithms like Go-Explore (Ecoffet et al., 2019), has the potential to simultaneously model and generative a diverse set of behaviors.




Figure 6. Histogram of steps to reach the goal node for random walks on the graph, shortest possible paths to the goal, and attempted shortest paths generated by the transformer model. ∞ indicates the goal was not reached during the trajectory.

Game	Random	Gamer
Breakout	2	30
Qbert	164	13455
Pong	-21	15
Seaquest	68	42055

Table 10. Atari baseline scores used for normalization.

Game	DT (Ours)	CQL	QR-DQN	REM	BC
Breakout	76.9 ± 27.3	61.1	6.8	4.5	40.9 ± 17.3
Qbert	2215.8 ± 1523.7	14012.0	156.0	160.1	2464.1 ± 1948.2
Pong	17.1 ± 2.9	19.3	-14.5	-20.8	9.7 ± 7.2
Seaquest	1129.3 ± 189.0	779.4	250.1	370.5	968.6 ± 133.8

Table 11. Raw scores for the 1% DQN-replay Atari dataset. We report the mean and variance across 3 seeds. Best mean scores are highlighted in bold. Decision Transformer (DT) performs comparably to CQL on 3 out of 4 games, and outperforms other baselines.

Game	DT (Ours)	10%BC	25%BC	40%BC	100%BC
Breakout	76.9 ± 27.3	10.0 ± 2.3	22.6 ± 1.8	32.3 ± 18.9	40.9 ± 17.3
Qbert	2215.8 ± 1523.7	1045 ± 232.0	2302.5 ± 1844.1	1674.1 ± 776.0	2464.1 ± 1948.2
Pong	17.1 ± 2.9	-20.3 ± 0.1	-16.2 ± 1.0	5.2 ± 4.8	9.7 ± 7.2
Seaquest	1129.3 ± 189.0	521.3 ± 103.0	549.3 ± 96.2	758 ± 169.1	968.6 ± 133.8

Table 12. %BC scores for Atari. We report the mean and variance across 3 seeds. Decision Transformer usually outperforms %BC.

Game	DT (Ours)	DT with no context ($K = 1$)
Breakout	267.5 ± 97.5	73.9 ± 10
Qbert	15.1 ± 11.4	13.6 ± 11.3
Pong	106.1 ± 8.1	2.5 ± 0.2
Seaquest	2.5 ± 0.4	0.6 ± 0.1

Table 13. Ablation on context length. Decision Transformer (DT) performs better when using a longer context length ($K = 50$ for Pong, $K = 30$ for others).