DECISION MAKING AS LANGUAGE GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Decision transformers are a recently proposed approach to offline reinforcement learning that leverages transformer-based auto-regressive sequence models. We discuss challenges associated with fine-tuning a given, pre-trained language model on a decision making task. We propose solutions to these challenges and study their viability on a shortest path problem. We also show how given language model allows us to bring to bear data-centric approaches to improving the model and how it opens up the possibility to treat the decision transformer objective as one task alongside others to perform transfer learning.

1 INTRODUCTION

Decision transformers (Chen et al., 2021; Janner et al., 2021) define the policy in an offline reinforcement learning task as a transformer-based, auto-regressive sequence model. They are trained supervised on sequences of triplets of actions, states and returns from roll-outs of an expert or other (including random) policy. Prompting with a high return at test-time encourages a decision transformer to follow action trajectories associated with high rewards and thereby to perform well.

Although decision transformers are a highly successful new approach to offline RL (cited several hundred times within few months of the publication of the basic approach by (Chen et al., 2021) and Janner et al. (2021)), applications of the model come with a significant shortcoming: the decision transformer is typically conceived of as an RL-specific architecture, with task-specific action-, state- and reward-heads, and is, as common in RL, trained on a narrowly defined set of target tasks (Zheng et al., 2022; Li et al., 2022). This is unlike foundation models in NLP, where the use of a single, canonical architecture applied to a wide range of tasks unlocks transfer learning and generalist models.

In this work, we study several improvements to the decision transformer based on treating the model as a pure language model. This makes it possible to pre-train and/or regularize the model with NLP tasks, while fine-tuning it on the task-specific RL objective. It also shifts the focus in RL tasks away from architecture search toward data-centric development. As an application of this data-centric view, we describe a novel approach to addressing the notorious problem of *reward-mismatch*: in many decision making tasks, the achievable return for an episode cannot be known ahead of generating the roll-out of the policy at run-time. This makes it impossible to know the appropriate return with which to prompt the model. We show that a simple data pre-processing step, that biases the model toward attempting to approximate the prompted return, allows us to achieve near-optimal performance by simply always prompting with the highest achievable reward.

2 RE-PURPOSING A PRE-TRAINED LANGUAGE MODEL AS DECISION TRANSFORMER

We follow Chen et al. (2021) and consider the decision making task as the task of modeling sequences of triplets (a, s, r), based on training data $\{(a_t^i, s_t^i, r_t^i)\}_{i=1}^N$ where a_t^i is an observed action at time t, s_t^i is an observed state at time t, and r_t^i is the cumulative return (or "return-to-go") that the given training trajectory incurred from time t onward. At inference time, the agent is prompted with a start state s and desired return r and generates an action a (followed by further states, returns and actions).

2.1 ENCODING

When using a pre-trained language model, it is necessary to commit to the model's tokenizer and vocabulary, and to express actions, states, returns as language. We restrict the encoding of these to ones that can be detected by a regular expression (as discussed below). Within these confines, we experimented with a variety of encoding schemes, including white-space separated values, XML-notation ("$\langle s \rangle 3 \langle / s \rangle \langle r \rangle 5 \langle / r \rangle \langle a \rangle \ldots$"), comma separated values, etc.

We found simple comma-separated notation ("$s = 3, r = 5, a = \ldots$") to work well and use that in all our experiments below. We note in passing that this encoding permits interleaving data from other tasks, including simple language, with the task-specific data.

2.2 ENVIRONMENT–MODEL INTERACTION AND SUB-STRING TRAINING

In this work, we model the environment as a simple Markov decision process (MDP). Since the goal of this work is to absorb the decision making task solely within a language modeling objective, the model does not have any task-specific outputs (or "heads"). In the absence of a task-specific output, it is necessary for the environment to continuously monitor and parse the model output to detect any well-formed actions within the model's output stream. We use a regular expression to this end, that matches on action outputs. While this provides sufficient flexibility in the encoding of actions (including all encoding schemes discussed above), it is simple and can admit highly efficient implementations.

Since we need to commit to the model's tokenizer, the output sequence generated by the model is not guaranteed to advance one-character at a time. For example, it is common to train auto-regressive language models using subword representations, such as byte-pair encoding (Sennrich et al., 2015), in which case a single model inference typically generates multiple characters. We let the environment monitor the model output by maintaining (i) a buffer of characters emitted by the model so far and (ii) a pointer to the end of the last well-formed action detected so far. Each emission by the model is then parsed with the regular expression, and the pointer is updated, if applicable.

Whenever a well-formed action is detected, the environment generates a state and reward, which are encoded and appended to the model's output stream. This approach of writing back into the model's buffer at run-time has been referred to as *environment forcing* by Recchia (2021), and it is used also by Nakano et al. (2021); Thoppilan et al. (2022); Cobbe et al. (2021) and others since.

Unfortunately, the character sequences at the transitions between model output and environment output are not guaranteed to contain token sequences that are common (or even contained at all) in the pre-training or fine-tuning dataset, and they depend on the encoding for actions, states and rewards. We observed in our experiments that this can prevent the pre-trained model from generating any reasonable decision sequences. This is a form of *exposure bias* (see, for example, Pang & He (2021): at inference time, the model sees token sequences that are rarely or never observed during training.

We found that the following, simple fix resolves this issue: for fine-tuning the model on decision sequences, we populate the training dataset with randomly generated crops of *proper sub-sequences* as a pre-processing step. The crops are generated on the character level. This ensures that (in the limit) each character occurs as the first or last character in a sequence and exposes the model to every character sequence that environment forcing may generate at inference time. We perform this step in all our experiments below and do no further processing nor any task-specific architecture adaptations to the model.

3 EXPERIMENTS

We study the viability of the language-based decision transformer on the shortest path problem. The shortest path problem (Wikipedia contributors, 2022) is a classic search problem, where the goal is to find the shortest path from a start node to a goal node in a given graph. It is common to use it as a reinforcement learning test bed by conditioning on a given, fixed graph (usually randomly generated) and presenting traversals of the graph alongside rewards. The goal of the agent is to reach the goal node through a sequence of transitions. Since the graph is not fully connected, not

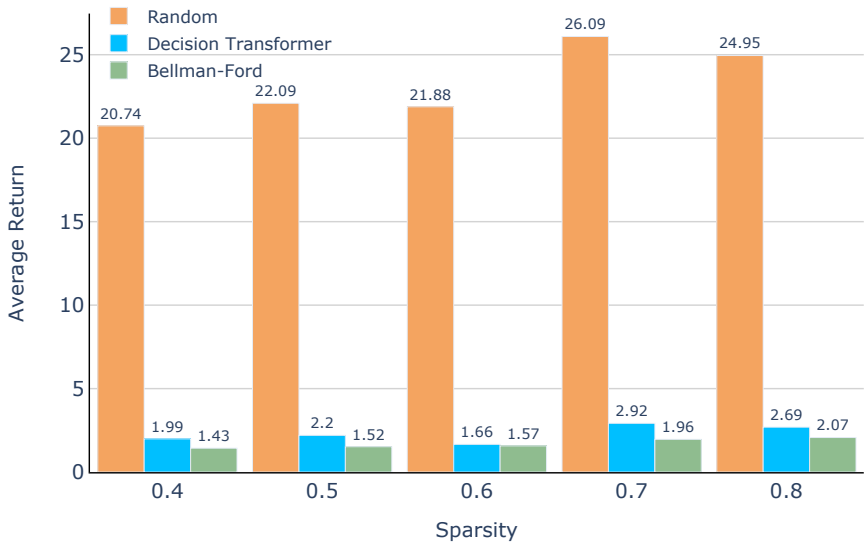

Figure 1: Performance of language-only decision transformer as compared to random walk and Bellman-Ford on a shortest path problem with varying levels of sparsity. The average number of steps to the goal are measured with 100 randomly generated 20-node graphs of each sparsity level.

all attempted transitions are valid. We adopt the reward structure proposed by Chen et al. (2021), which amounts to generating a reward of -1 for each attempted transition, including those that are invalid, and generating a reward of 0 after the goal state is reached.

Following Chen et al. (2021), we train the model on random traversals (of length 30 steps) of randomly generated graphs with 20 nodes and varying levels of connectivity. We use $60k$ sequences for training and train for a single epoch. We use the data pre-processing step discussed in the next section, and at inference time prompt the model with a random initial state and an optimal reward request of -1. The model interacts with the environment using regular expression parsing as discussed above. Figure 1 compares the performance of the language-based decision transformer ("DT" in the figure) with the performance of a random walk ("Random") and the optimal performance obtained using Bellman-Ford ("Bellman-Ford").

3.1 Turning reward prompts into reward requests

Reward conditioning in a decision transformer amounts to "selecting" from the trajectories learned by the model one which is associated with high return. When a conditioned reward is not achievable given the state, this "selection" fails and can generate behaviours not encountered during training and thereby generate trajectories with very bad performance (Chen et al., 2021; Zheng et al., 2022).

Unfortunately, conditioning on an achievable return is not possible at inference time because it is not usually known which return is achievable for any given state. Chen et al. (2021) propose resolving this issue by incorporating a prior into the sampling procedure used by the model.

In this work, we explore a simple and entirely data-centric alternative, which does not amount to modifying the model architecture or sampling mechanism. Our approach amounts to pre-processing the training data used for fine-tuning the model to encourage the model to generate the best possible approximation for a given return. It amounts to effectively changing the semantics of return prompts from "trajectory selection" to "return request". To this end, after generating the initial training data, we replace the *first* return encoding to a random integer (which essentially plays the role of

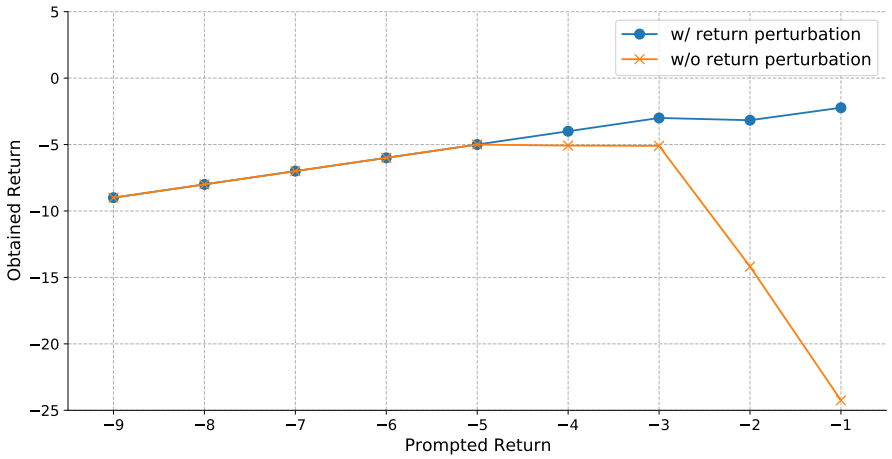

Figure 2: Comparison of obtained return (y-axis) for varying return prompts (x-axis) for a model trained without perturbed reward requests ("w/o return perturbation") and a model trained with perturbed reward requests ("w/ return perturbation").

a "return request") r with a random perturbation $\hat{r}$ such that $\hat{r} < r$. The rest of the sequence remains untouched. This results in an overly optimistic return request, which the model learns to replace with a more realistic return estimate in the subsequent time-step. This allows the model to learn to replace overly demanding return requests by the next best *achievable* return estimate and to subsequently follow the resulting trajectory. We perform this pre-processing step before the random crop generation discussed above.

Figure 2 shows the result. For state/return values that are achievable by the model, conditioned and obtained returns are highly correlated, but for the vanilla decision transformer, obtained returns degrade sharply for overly large return prompts. Performing the data pre-processing step before training (labelled "w/ return perturbation"in the figure) changes this behaviour and allows us to simply prompt with the best achievable return after training.

4 Discussion

We discussed the challenges associated with turning a pre-trained, *fixed* language model into a decision transformer, as well as solution to those challenges. We demonstrated the viability of these on a shortest path problem. While our results show both the possibility and advantages of using pre-trained language models for decision making (for example, by allowing for flexible return prompts), in future work we aim to study these in more challenging tasks. Pre-trained language models may come with additional advantages, such as the ability to mix in with other, RL or non-RL, objectives. This may greatly benefit the current push towards more generalist models (for example, Reed et al. (2022)). It would also be interesting to explore data-driven (as opposed to architectural) approaches to deal with environment stochasticity (Paster et al., 2022).

Another avenue for future research is to mix action emissions with other textual outputs. The interactions with the environment via regular expressions would allow the model to interleave language-based reasoning with the emission of actions. In complex tasks, it would be possible for the model to allocate varying numbers of tokens (and hence, computation-time) to the task, making it possible to study the benefits of adaptive computation time (Graves, 2016) in the context of modern foundation models.

REFERENCES

Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling, 2021.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021.

Alex Graves. Adaptive computation time for recurrent neural networks, 2016. URL https://arxiv.org/abs/1603.08983.

Michael Janner, Qiyang Li, and Sergey Levine. Reinforcement learning as one big sequence modeling problem, 2021.

Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, Jacob Andreas, Igor Mordatch, Antonio Torralba, and Yuke Zhu. Pre-trained language models for interactive decision-making, 2022.

Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. Webgpt: Browser-assisted question-answering with human feedback, 2021. URL https://arxiv.org/abs/2112.09332.

Richard Yuanzhe Pang and He He. Text generation by learning from demonstrations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=RovX-uQ1Hua.

Keiran Paster, Sheila McIlraith, and Jimmy Ba. You can't count on luck: Why decision transformers fail in stochastic environments, 2022.

Gabriel Recchia. Teaching autoregressive language models complex tasks by demonstration. *ArXiv*, abs/2109.02102, 2021.

Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent, 2022. URL https://arxiv.org/abs/2205.06175.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. Lamda: Language models for dialog applications, 2022.

Wikipedia contributors. Shortest path problem — Wikipedia, the free encyclopedia, 2022. URL https://en.wikipedia.org/w/index.php?title=Shortest_path_problem&oldid=1109389165. [Online; accessed 24-September-2022].

Qinqing Zheng, Amy Zhang, and Aditya Grover. Online decision transformer, 2022.