Masked Autoencoding for Scalable and Generalizable Decision Making

Anonymous Author(s)
Affiliation
Address
email

Abstract

We are interested in learning scalable agents for reinforcement learning that can learn from large-scale, diverse-quality sequential data similar to current large vision and language models. To this end, this paper presents masked decision prediction (MaskDP), a simple and scalable self-supervised pretraining method for reinforcement learning (RL) and behavioral cloning (BC). In our MaskDP approach, we employ a masked autoencoder (MAE) to state-action trajectories, wherein we randomly mask state and action tokens and reconstruct the missing data. By doing so, the model is required to infer masked out states and actions and extract information about dynamics. We find that masking different proportions of the input sequence significantly helps with learning a better model that generalizes well to multiple downstream tasks. In our empirical study we find that a MaskDP model gains the capability of zero-shot transfer to new BC tasks, such as single and multiple goal reaching, and it can zero-shot infer skills from a few example transitions. In addition, MaskDP transfers well to offline RL and shows promising scaling behavior w.r.t. to model size. It is amenable to data efficient finetuning, achieving competitive results with prior methods based on autoregressive pre-training.

1 Introduction

Self-supervised pre-training has made tremendous successes in natural language processing (NLP) and vision [11, 8, 3, 4]. These methods work by predicting a removed portion of the data, which is often referred to as masked token prediction. Masked token prediction enables highly successful methods for pre-training in NLP and vision, $e.g.$, Transformer [27], GPT [4], BERT [8], and MAE [11]. There has been a lot of evidence indicating that these pre-trained representations generalize well to various downstream tasks. The cornerstone of these successes is masked token prediction, a simple and scalable self-supervised pre-training objective, that can excellently leverage unlabeled data which is widely available on internet and easy to collect [8, 11, 1].

The idea of masked token prediction is natural and applicable in reinforcement learning (RL) as well. Driven by the successes in NLP and vision, prior work has explored using masked token prediction within offline RL $e.g.$ decision transformer (DT) [5] and trajectory transformer (TT) [12]. These methods are based on autoregressive next token prediction, similar to GPT [4]. However, these works show no sign of generalizing to downstream tasks or leveraging unlabeled diverse data. In addition, DT [5] needs reward-labeled expert-level data, while TT [12] requires discretizing states and actions, further limiting its applicability. Despite significant interest, progress of masked token prediction in RL significantly lag behind NLP and vision.

In this work, we propose a method named Masked Decision Prediction (MaskDP) to learn generalizable models that achieve data-efficient adaptation to various downstream tasks. MaskDP is a

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

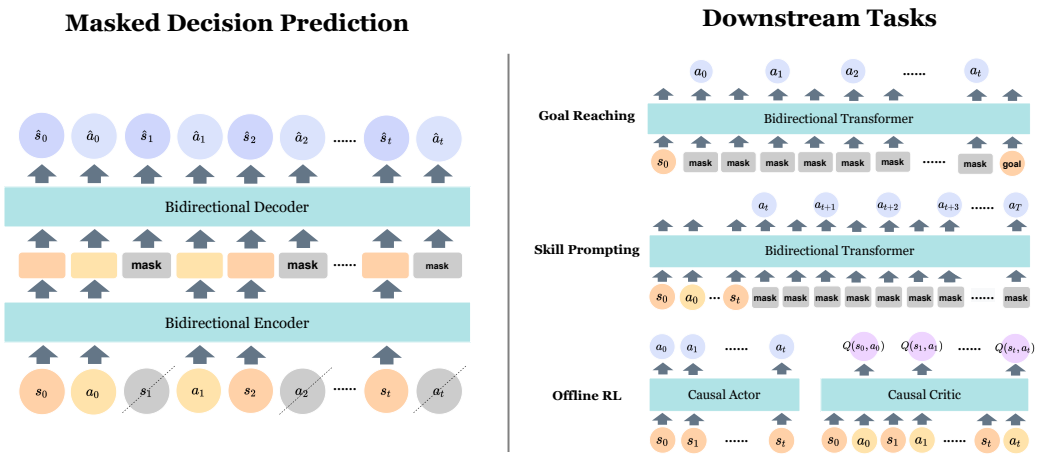

Figure 1: Illustration of MaskDP. During pretraining stage, we perform the masked token prediction task. And after pretraining, the model can be deployed to various downstream tasks using different mask patterns.

self-supervised pre-training method that can leverage unlabeled diverse data. With MaskDP pre-training, the model can generalize well to both goal reaching and offline RL, two distinctive and popular RL paradigms.

Our first key observation is that masked token prediction with random masking similar to BERT [8] and MAE [11] provides a general and flexible way of fitting unsupervised data. Unlike autoregressive next token prediction used in prior works, random masking requires the model to infer masked out states and actions, leading the model to extract more information about forward and inverse dynamics from each sample.

Our second key observation is that since states and actions are highly correlated temporally, trajectories have significantly lower information density, $i.e$, it is easier to predict action or state based on nearby states and actions. Consequently, a high mask ratio (95%) is necessary to make reconstruction task meaningful. Unlike in MAE [11] and BERT [8] where the goal is learning representations, we want to directly apply MaskDP to various downstream tasks, and ifferent mask ratios induce different pre-train and downstream gaps. For example, consider the goal-reaching task within certain time limit. Given current state, future goal and mask tokens between them, the model should be able to inpaint intermediate actions as the goal-reaching plan. The mask ratio varies from short-term plans to long-term plans. Therefore, we combine multiple different mask ratios ($e.g.$ 15%, 35%, 75%, and 95%), and mask a portion of data using a randomly sampled mask ratio. Our experiments show that doing so is crucial to achieving high performance. We show that self-supervised pre-trained MaskDP achieves high performance in challenging multiple goals reaching setting, outperforming strong baselines in a zero-shot manner. MaskDP is also amenable to fine-tuning–we find that finetuned MaskDP learns significantly faster than training from scratch or pretrained baselines in offline RL.

We highlight our key results here:

- **Single goal reaching**: MaskDP achieves performance that exceeds or matches both training from scratch task-specific methods and other pre-training based methods.
- **Sequential multiple goal reaching**: MaskDP can reach a sequence of goals effectively, even without closed-loop execution, while outperforming iterative baselines significantly.
- **Offline RL**: MaskDP achieves competitive results as specialized approaches. Notably, the strongest alternative models are GPT-based pre-trained models which naturally fit the setting for RL tasks, whereas our architecture is generally applicable across tasks.

2 Related work

Masked modeling in language and vision. Large-scale language models are highly successful [8, 4]–after pretraining on a large amount of data, these pre-trained representations generalize well to

various downstream tasks. Taking inspiration from the success in NLP, Transformer [27] based methods have been proposed to model images [6, 9, 3, 11]. iGPT [6] operates on sequences of pixels and predicts unknown pixels. BEiT [3] proposes to predict discrete tokens [26, 18]. MAE [11] proposes to randomly mask patches of the input image and reconstruct the missing pixels. Since we apply random mask across states and actions, our work is also related to prior work on masked prediction across multiple input modalities [see e.g. 28].

Masked modeling in RL Masking trajectories to discover learning signal has been studied in RL [17, 19, 5, 12, 33]. MVP [29] studies transferring pre-trained visual representations to RL tasks. Modelling inverse dynamics has also been studied for robot learning from demonstrations and sim-to-real transfer [7, 25]. TT [12] studies autoregressive next token prediction for model-based RL applications. DT [5, 33] study masking autoregressive next token prediction conditioned on return. ICM [17] and SPR [19] study predicting masked state and action in a transition tuple for exploration.

Different from these works, MaskDP randomly masks a portion of trajectories and generalizes prior masking strategies such as inverse dynamics. In addition, MaskDP generalizes well to downstream tasks while prior work is task-specific.

Unsupervised pretraining in RL Our work falls under the category of self-supervised pretraining in RL. Self-supervised discovery of a set of task-agnostic behaviors by means of seeking to maximize an intrinsic reward has been explored as intrinsic motivation [2], often with the goal of encouraging exploration [22, 16]. APT [15] studies nonparametric entropy maximization for pretraining and is extended to learning skills [13, 14]. Proto-RL [30] further improves pretraining. APV [21] shows successful transfer of pretrained representation across domains. Many of these methods are used to pretrain agents that are later adapted to specific reinforcement learning tasks. Using offline data for pretraining agents has also been explored in prior work [20, 23, 32]. SGI [20] proposes combining self-predictive representation [19] and inverse dynamics prediction. ATC [23] studies contrastive pretraining on trajectories.

3 Method

MaskDP masks and reconstructs state-action sequences during the pretraining stage. After that, it can be zero-shot deployed or finetuned for various downstream tasks. The paradigm of the model for pretraining and finetuning is summarized in Figure 1.

3.1 MaskDP Pretraining

Random masking. For sequences with low information density, a *high* masking ratio is typically applied to eliminate information redundancy and make the task sufficiently difficult to avoid trivial interpolation from visible neighbor tokens. However, unlike vision and language, where the goal is to learn good representations; we also consider directly deploying this model by leveraging its inpainting capability for various downstream tasks. For example, we can give the model a goal at timestep T and mask all the future inputs, the model can generate intermediate actions by inpainting the mask tokens. The mask ratio varies from goals to goals, depending on the time budgets. To reduce the gap between training and deployment, we keep a set of mask ratios ($i.e.$ 15%, 35%, 50%, 75%, and 95%), and the data is randomly masked with a ratio sampled from this set. We find that masking multiple proportions of the input yields a meaningful self-supervisory task.

We apply random masking on state tokens and action tokens independently. By doing so, the model is implicitly learning both forward and inverse dynamics. This also provide more flexibility as we can provide state or action-level inputs but not transition-level.

Architecture Our encoder is a Transformer [27] but applied only on visible, unmasked states and actions, similar to MAE [11]. The states and actions are first embeded by separated linear layers, positional embeddings are then added, and lastly the embeddings are processed by a series of self-attentional blocks. The decoder operates on the full set of encoded visible state and action tokens and mask tokens. Each mask token is a shared, learned vector that indicates the presence of a missing token to be predicted. Similarly to the encoder, the masked whole sequence will pass through

separated linear projections added with positional embedding prior to being passed to the decoder. Both the encoder and decoder are bidirectional.

Prediction target Our MaskDP reconstructs the input by predicting the whole action and state sequences. The last layer of the decoder consists of two MLPs to decode states and actions separately. The loss function computes the mean squared error (MSE) between the reconstructed whole sequence and original inputs. Different from other masked prediction variants [11, 8], we found mask loss is not useful in our setting, as our goal is to obtain an scalable decision making model but not only for representation learning.

3.2 MaskDP Downstream Tasks

MaskDP for goal reaching We consider the problem of reaching one goal or multiple goals from a given state. The model has to generate a sequence of actions to reach goals within a certain amount of steps. MaskDP denoising pretraining objective fits the goal reaching scenario well as the model must learn to inpaint masked actions based on remaining states. In this task, the MaskDP encoder input is a concatenation of initial state and goals, and the decoder input is a concatenation of initial state embedding, masked token sequence, and goal embeddings. Note that the number of masked tokens determines the number of timesteps the model is expected to reach the given goals. The model then generates a state-action sequence, where we can directly execute the whole action sequence (namely "open-loop"), or only execute the first action and forward the model again with the obtained new observation (namely "closed-loop").

MaskDP for skill prompting Skill prompting requires the model to generate a trajectory in the form of some given context. For example, consider a walker agent: if we prompt it with states and actions of walking/running/standing, it should continue to generate a trajectory in the same skill pattern. This requires the model to fill the future masked sequence based on the initial inputs. Specifically, we append future mask tokens to the initial state-action sequences. The model can also forward once to generate the whole future, or refill the mask tokens every time. To make it consistent with goal-reaching task, we still refer them as "open-loop" and "closed-loop" respectively.

MaskDP for offline RL In offline RL, the objective is to learn one model for maximizing return for a task specified by a reward function. This is different from our self-supervised pretraining target, so extra finetuning is needed. We adopt standard actor-critic framework similar to TD3 [10] by adding a critic head and actor head, where the actor takes a state sequence as input, and the critic takes the state-action sequence as input. Both are mask-free. To match the setting in RL, we change the bidirectional attention mask in the transformer to a causal attention mask. More details about RL finetuning can be found in section 4.2.3.

4 Experiments

In our experiments, we evaluate transfer learning in downstream tasks using MaskDP. Section 4.1 introduces the environments, pre-training, and the baselines compared in experiments. Section 4.2 summarizes the results of MaskDP on goal reaching, skill prompting, and offline RL. Through further analysis in Section 4.3, we present an ablation study on various design choices of our model.

4.1 Experiments Setup

Environments: domains vs. tasks We adopt the environment setup used in EXoRL [31], based on DeepMind control suite [24], where a domain describe the type of agent (e.g. Walker) but tasks are specified by rewards (e.g., Walker walk, Walker run). We use 3 domains (Walker, Cheetah and Quadruped) with 7 tasks in total. More details about the environments can be found in the Appendix.

Pre-training datasets Real-world pretraining data generally varies greatly in quality. To mimic this, we construct two different pretraining datasets to approximate different data quality scenarios.

- **Near-expert**: For every task, we train a TD3 agent [10] for 1M steps and freeze its parameters. We rollout the policy with Gaussian random noise and collect 4K trajectories on each task.

- **Mixed**: This dataset consists of diverse data collected from various agents, including 2K near-expert trajectories for each task. Similar to ExoRL [31], we collect 10K exploratory trajectories using intrinsic reward from Proto-RL [30] for each domain. We also to use a TD3 [10] agent to maximize the sum of extrinsic reward and the Proto-RL intrinsic reward, and store its 2k experience on each task.

For more details about the above datasets, please refer to Section A. We perform both single-task and multi-task pretraining using the above datasets. The former leverages task-specific data while the latter utilizes data from all tasks within the same domain. We pretrain agents for 400k gradient steps. Specifically, for the model pretrained on the near-expert dataset, we perform zero-shot evaluation of goal reaching and skill prompting, and finetuning for offline RL; for model trained on the mixed dataset, we provide the finetuning results in Section 4.3 and Section A.

Baselines

- GPT. We train an autoregressive model similar to GPT [4] which takes the past states and actions as input to predict the next state or action.
- Goal-GPT. We specifically modify GPT to Goal-GPT to evaluate its performance on goal reaching tasks. The model takes current goal and observations as input, and predicts the action to reach this goal. The model is trained using a behaviour cloning loss as [5].
- Goal-MLP. Standard behavior cloning method that conditions on the goal. The major difference between this and Goal-GPT is here we do not use the causal Transformer architecture to make the history visible.

By default, MaskDP uses a 3-layer encoder and 2-layer decoder, and the baselines based on GPT use 5 attention layers. MaskDP and all the above models are comparable with similar architecture design and size, and share the same training hyper-parameters. Details about the architecture and training of MaskDP and the above baselines can be found in Section A.

4.2 Main Results

4.2.1 Goal reaching

We consider both single and multiple goal-reaching settings. The agent is required to reach one or multiple goals from an given state, which are all sampled from the same trajectory to guarantee reachabilty within a reasonable time budget. During evaluation, the agent rolls out to reach the given goal(s) within a time budget. The evaluation dataset is also collected by the same RL agent in 3 environments with different seeds, which is unseen during pre-training. The detailed settings are:

- **single-goal reaching**: For every trajectory in the validation set, we randomly sample a start state and a future state in $T \in [15, 20)$ steps as the goal. All the methods are evaluated on the same set of 300 state-goal pairs with a given budget of $T + 3$. We set the agent to the start state[1] and report the L2 distance between the goal and the closest rollout state within this budget.
- **multi-goal reaching**: For every trajectory in the validation set, we randomly sample a start state and 5 goal states at random future timesteps from $[12, 60)$. We evaluate the same set of 100 state-goal sequences and adding additional 5 timestep budgets for all the goals. Similar to single-goal reaching, We report the L2 distance between every goal and the closest rollout state before running out of its corresponding budget.

We show the zero-shot performance of MaskDP and baselines pretrained with the near-expert data (both in single-task and multi-task setting). We report L2 distance averaged over the states and goals sampled based on the above rules. Tables of run numbers and standard derivations can be found in Section B.

Figure 2 show the results of reaching single goal. The y-axis is the L2 distance (the lower the better). We observe that both MaskDP (open-loop) and MaskDP (closed-loop) outperform Goal-GPT and Goal-MLP. Despite Goal-GPT being a natural formulation for goal reaching, MaskDP reaches a lower distance to goal. We attribute the effectiveness to learning a better understanding of the forward

[1] We accomplish this by resetting the simulator.

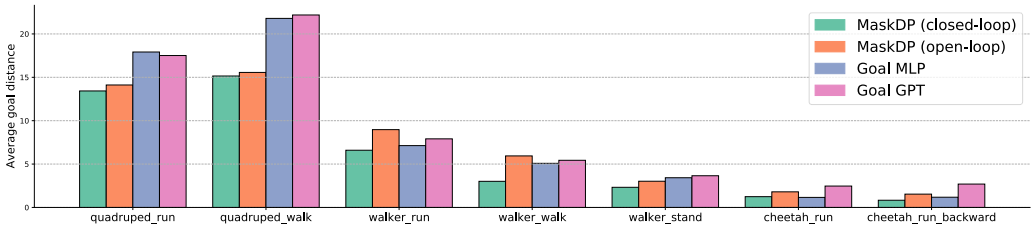

Figure 2: Single task pretraining followed by single goal reaching downstream task. MaskDP with closed-loop execution achieves the best performance on all the tasks, and get the most significant improvements in the Quadruped domain, which is higher dimensional.

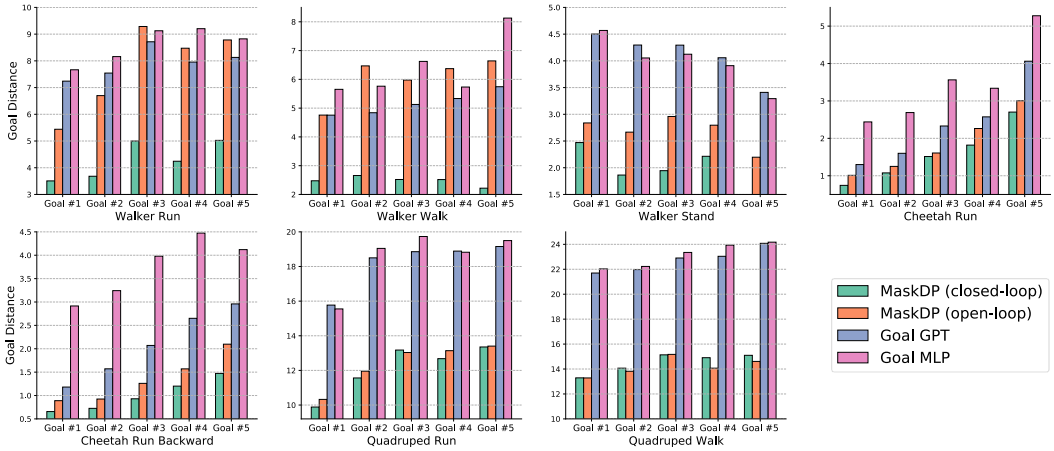

Figure 3: Single task pretraining followed by multiple goals reaching downstream task. MaskDP achieves significant improvement on all the tasks with better flexibility in sequential goal reaching.

and inverse dynamics implicitly. We also observe that the advantages of MaskDP are even more significant on higher dimensional environments, such as Quadruped.

For the more challenging multi-goal reaching task, MaskDP has a significant advantage in flexibility: we can just provide the goals at specific time budgets with interleaved masks and get an executable plan; however, for Goal-MLP and Goal-GPT, we have to change goals at certain timesteps to fulfill future multiple goals. As shown in Figure 3, MaskDP outperforms both goal-GPT and BC by a large margin. We hypothesize that this is due to having "foresight" about future goals, which can help the agent to produce a better plan.

We can get similar conclusions from the multi-task pretrained models in Figure 4 and Figure 5, where our method consistently works well on all domains, with the most visible advantage in multi-goal reaching setup.

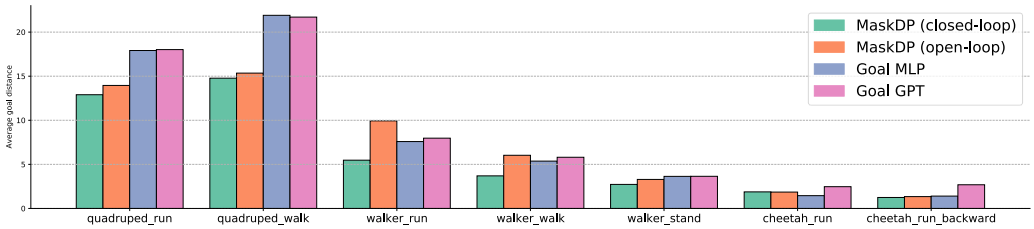

Figure 4: Multiple tasks pretraining followed by single goal reaching downstream task, where MaskDP with closed-loop execution works the best, especially in the Quadruped domain.

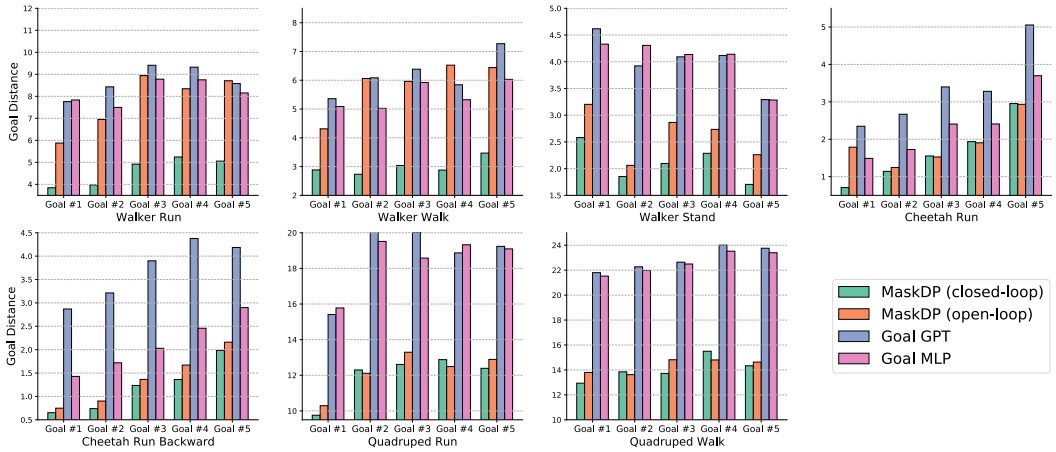

Figure 5: Multiple task pretraining followed by multiple goals reaching downstream task.

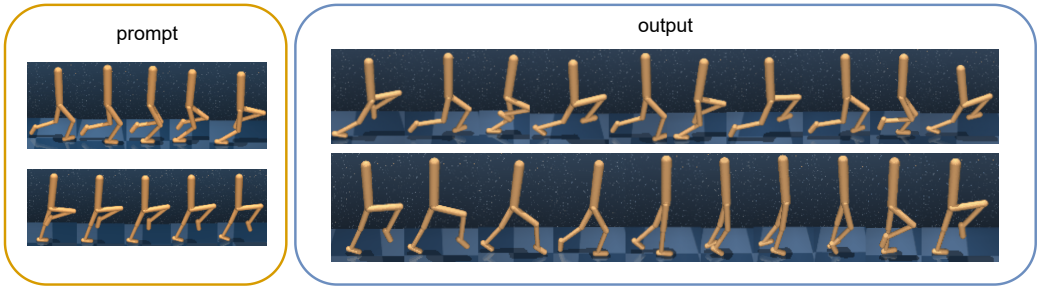

Figure 6: Qualitative results on for skill prompting in the Walker domain. Given 5 initial states, the model learns to forecast future trajectories as in the expert-level behaviour.

4.2.2 Skill Prompting

We are interested in the learned behavior of pretrained models. We use prompting, which has become popular in analyzing models ever since GPT [4]. To do so, we give the agent a short state-action segment randomly cropped from an expert trajectory, set the agent to the last state of the segment, and let the model continue to generate consecutive behaviors. We evaluate the quality of the generated sequence by comparing its obtained rewards with the rollout of a skilled expert.

To be specific, we prompt the multi-task model trained with expert data and sample a 5-timestep state-action segment from $T \in [100, 900]$, where the agent can be walking/running at low or high speeds. We prompt the model with this short segment and let the model generate rollouts for 20, 40, 60 timesteps. We provide both qualitative results and quantitative results in Figure 7 and Figure 6 respectively. We can see in Figure 7, both our method and GPT can match the expert return. Although our method is not trained in an autoregressive way, it can still perform well in the sequence generation task. For the zero-shot prompting results of MaskDP/GPT trained on mixed data, see Section B.

4.2.3 Offline Reinforcement Learning

Evaluation We provide a 2M buffer of the data collected by Proto-RL [30] as in ExoRL [31] does, where the overall return of the data is quite low and thus the BC-based method cannot work well. ExoRL [31] simply shows that an offline TD3 agent works the best on diverse low-return offline data.

We can modify MaskDP to this setting by adding additional actor and critic heads and perform TD learning. We evaluate the efficiency of the pretrained model by its return after certain TD gradient steps. The results are shown in Figure 8 averaged over 3 seeds. We observe MaskDP is capable of adapting to downstream tasks quickly, outperforms training from scratch, and achieves similar results as the GPT baseline. Note that in this setting, we need to replace the bidirectional attention mask with causal attention mask, so there is a larger gap between pre-training and downstream tasks finetuning

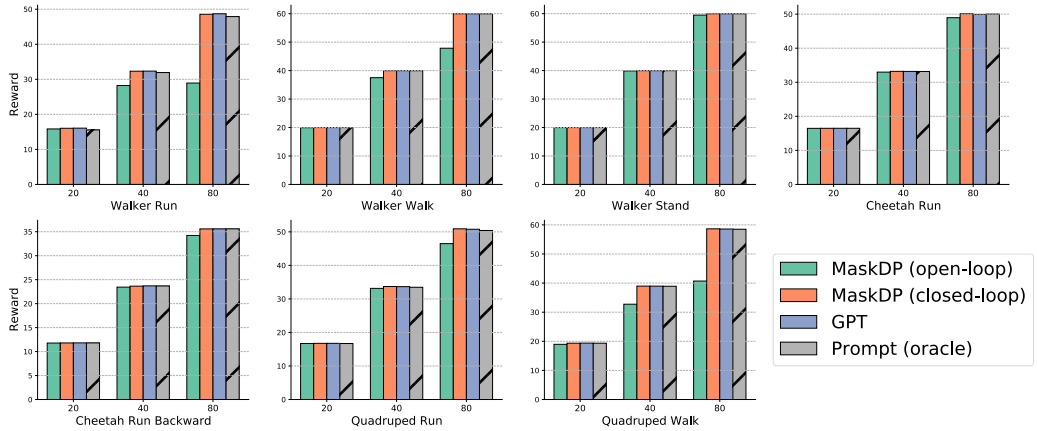

Figure 7: Quantitative results on learned behaviors using prompt. Both MaskDP and GPT can match or even slightly surpass the expert-level performance (right grey bar) in trajectory forecasting.

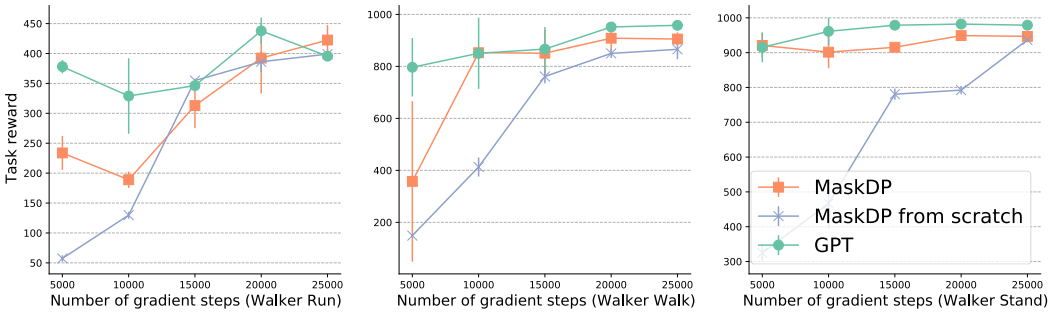

Figure 8: Offline RL results on Walker domain. The result of MaskDP matches the GPT-style pretraining performance, and are all comparable to the SoTA in ExoRL [31].

compared with GPT, which is trained with causal masking. Note that MaskDP from scratch is almost the same as GPT from scratch (both with causal masking). From Figure 8, both MaskDP and GPT can match the best result in ExoRL [31] from their offline TD3 agent, where BC-based method cannot successfully solve this task.

4.3 Analysis

Model scalability We also pretrain our agent and baselines on the diverse "mixed" dataset. We compare smaller version of MaskDP and Goal-GPT[2] on the Quadruped goal reaching problem as shown in Figure 9. The x-axis is the finetuning gradient steps on the expert dataset, and y-axis in the L2 distance to the goal (the lower the better). We found for both zero-shot evaluation and finetuning, our model's performance improves when model size is enlarged, whereas for Goal-GPT the performance gain is not obvious.

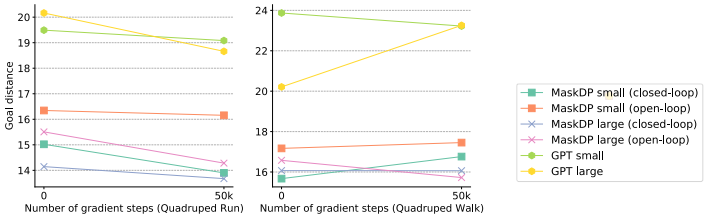

Figure 9: Model scalability.

Mask ratio ablation Figure 10 shows the influence of the masking ratio. With a fixed mask ratio, we observe that an extremely high mask ratio (95%) generally does not work well and the typical

[2]Both are in 3 attention layers.

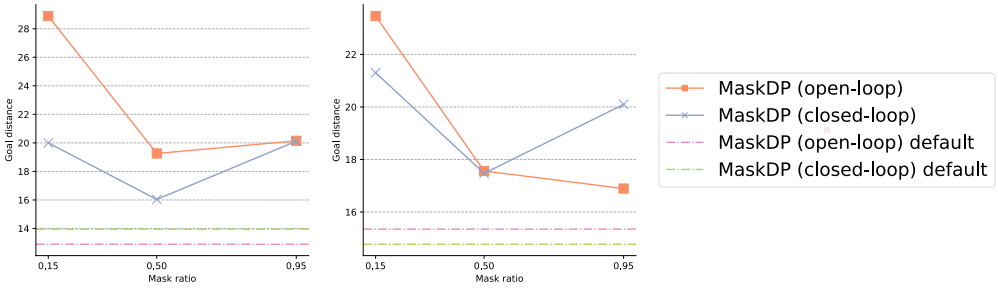

Figure 10: Mask ratio ablation. We compare our multiple ratio pretrained model with models trained with fixed ratios, where our masking strategy can achieve much better performance.

mask ratio (15%) used in BERT seems to perform much worse than others. A middle mask ratio 50% performs reasonably well, despite still being surprisingly high, similiar to the observations in MAE. However, our mixed mask ratio strategy strictly outperform all the above options.

Predicting unmasked tokens ablation We also compare the model trained with mask loss vs. total loss. As shown in Figure 11, empirically we do not find mask loss has more advantage than total loss, even on the relatively clean expert dataset, it converges slower than using total loss. For the results on diverse data, please refer to Section B.

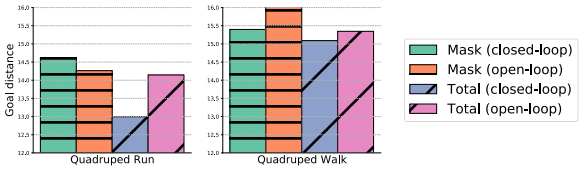

Figure 11: Masked loss and total loss ablation on 20k pretraining gradient steps. The model trained with total loss converges faster than the one trained with masked loss.

5 Conclusion

This paper presents masked decision prediction (MaskDP), a simple and scalable self-supervised method for reinforcement learning (RL) inspired by current large language and vision models. MaskDP is capable of learning scalable and generalizable agents for reinforcement learning that can learn from diverse-quality data sources and infer tasks in goal-reaching and skill-execution settings. Through our empirical study we find MaskDP models ourperform past work in zero-shot goal reaching and transfer well to downstream RL tasks, performing competitively with prior pre-training and training from scratch methods.

5.1 Limitations and Future Work

CV and NLP domains have shown that the true promise of masking architectures lies with their ability to ingest diverse, fully unsupervised data. In the future, we will investigate how MaskDP performs when trained without access to any expert data and with data from distributions and tasks far different from the downstream task.

The architecture used in MaskDP closely resembles a model-based method, as states are predicted sequentially from actions. In this paper, we use use the predicted next actions directly as this is the simplest and fastest approach. However, we could easily extend MaskDP to plan through our learned model and compare against related baselines such as [12].

5.2 Societal Impact

This is an algorithm for training agents in the style of recent large-scale CV and NLP models. While we do not anticipate particular social risks from our method, as algorithms become capable of ingesting large-scale, in-the-wild data it is important to ensure the dataset does not reinforce undesirable biases or promote harmful behaviors.

References

[1] J.-B. Alayrac, J. Donahue, P. Luc, A. Miech, I. Barr, Y. Hasson, K. Lenc, A. Mensch, K. Millican, M. Reynolds, R. Ring, E. Rutherford, S. Cabi, T. Han, Z. Gong, S. Samangooei, M. Monteiro, J. Menick, S. Borgeaud, A. Brock, A. Nematzadeh, S. Sharifzadeh, M. Binkowski, R. Barreira, O. Vinyals, A. Zisserman, and K. Simonyan. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv: Arxiv-2204.14198*, 2022.

[2] G. Baldassarre and M. Mirolli. *Intrinsically motivated learning in natural and artificial systems*. Springer, 2013.

[3] H. Bao, L. Dong, and F. Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.

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

[5] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *arXiv preprint arXiv: Arxiv-2106.01345*, 2021.

[6] M. Chen, A. Radford, R. Child, J. Wu, H. Jun, D. Luan, and I. Sutskever. Generative pretraining from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020.

[7] P. Christiano, Z. Shah, I. Mordatch, J. Schneider, T. Blackwell, J. Tobin, P. Abbeel, and W. Zaremba. Transfer from simulation to real world through learning deep inverse dynamics model. *arXiv preprint arXiv:1610.03518*, 2016.

[8] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[9] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[10] S. Fujimoto, H. van Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv: Arxiv-1802.09477*, 2018.

[11] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv:2111.06377*, 2021.

[12] M. Janner, Q. Li, and S. Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34, 2021.

[13] M. Laskin, H. Liu, X. B. Peng, D. Yarats, A. Rajeswaran, and P. Abbeel. Cic: Contrastive intrinsic control for unsupervised skill discovery. *arXiv preprint arXiv:2202.00161*, 2022.

[14] H. Liu and P. Abbeel. Aps: Active pretraining with successor features. In *International Conference on Machine Learning*, pages 6736–6747. PMLR, 2021.

[15] H. Liu and P. Abbeel. Behavior from the void: Unsupervised active pre-training. *Advances in Neural Information Processing Systems*, 34, 2021.

[16] P.-Y. Oudeyer and F. Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.

[17] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR, 2017.

[18] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.

[19] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. Courville, and P. Bachman. Data-efficient reinforcement learning with self-predictive representations. *arXiv preprint arXiv:2007.05929*, 2020.

[20] M. Schwarzer, N. Rajkumar, M. Noukhovitch, A. Anand, L. Charlin, R. D. Hjelm, P. Bachman, and A. C. Courville. Pretraining representations for data-efficient reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[21] Y. Seo, K. Lee, S. James, and P. Abbeel. Reinforcement learning with action-free pre-training from videos. *arXiv preprint arXiv:2203.13880*, 2022.

[22] Ö. Şimşek and A. G. Barto. An intrinsic reward mechanism for efficient exploration. In *Proceedings of the 23rd international conference on Machine learning*, pages 833–840, 2006.

[23] A. Stooke, K. Lee, P. Abbeel, and M. Laskin. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, pages 9870–9879. PMLR, 2021.

[24] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. de Las Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, T. Lillicrap, and M. Riedmiller. Deepmind control suite. *arXiv preprint arXiv: Arxiv-1801.00690*, 2018.

[25] F. Torabi, G. Warnell, and P. Stone. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018.

[26] A. Van Den Oord, O. Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

[27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[28] Y. Wang, S. Joty, M. R. Lyu, I. King, C. Xiong, and S. C. Hoi. Vd-bert: A unified vision and dialog transformer with bert. *arXiv preprint arXiv:2004.13278*, 2020.

[29] T. Xiao, I. Radosavovic, T. Darrell, and J. Malik. Masked visual pre-training for motor control. *arXiv preprint arXiv:2203.06173*, 2022.

[30] D. Yarats, R. Fergus, A. Lazaric, and L. Pinto. Reinforcement learning with prototypical representations. In *International Conference on Machine Learning*, pages 11920–11931. PMLR, 2021.

[31] D. Yarats, D. Brandfonbrener, H. Liu, M. Laskin, P. Abbeel, A. Lazaric, and L. Pinto. Don't change the algorithm, change the data: Exploratory data for offline reinforcement learning. *arXiv preprint arXiv: Arxiv-2201.13425*, 2022.

[32] A. Zhan, P. Zhao, L. Pinto, P. Abbeel, and M. Laskin. A framework for efficient robotic manipulation. *arXiv preprint arXiv:2012.07975*, 2020.

[33] Q. Zheng, A. Zhang, and A. Grover. Online decision transformer. *arXiv preprint arXiv: Arxiv-2202.05607*, 2022.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes]
- Did you include the license to the code and datasets? [Yes]
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Will release the code in the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We provide comprehensive details about our experiments in the Appendix A.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See our experiments in section 4.2.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]