

Appendix: Shortest-Path Constrained Reinforcement Learning for Sparse Reward Tasks

A. More ablation study




Figure 9. Average episode reward of **SPRL** with varying $k = 1, 3, 10, 30$ as a function of environment steps for *DeepMind Lab* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained with $k = 10$.




Figure 10. Average episode reward of **SPRL** with varying $k = 3, 5, 7$ as a function of environment steps for *MiniGrid* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained with $k = 5$.

A.1. Effect of k

As proven in **Lemma 2** and shown in Section 6.6, the larger k , the k -shortest constraint promises a larger reduction in policy space, which results in faster learning. However, with our practical implementation of **SPRL** with a learned (imperfect) reachability network, overly large k has a drawback. Intuitively speaking, it is harder for policy to satisfy the k -shortest constraint, and the supervision signal given by our cost function becomes sparser (*i.e.*, almost always penalized). Figure 9 and 10 shows the performance of **SPRL** on *DeepMind Lab* and *MiniGrid* domains with varying k . In both domains, we can see that there exists a “sweet spot” that balances between the reduction in policy space and sparsity of the supervision (*e.g.*, $k = 10$ for *DeepMind Lab* and $k = 5$ for *MiniGrid*). In practice, we performed grid-search over the hyper-parameter k .

Training multiple reachability networks with different k ’s may achieve the advantages of both low and high k . But practically, training multiple reachability networks requires high computational cost and more extensive hyperparameter search, which can be intractable. Thus, we instead tried the curriculum learning of k ; *i.e.*, starting from low k and gradually increasing it up to the target k . However, we found that curriculum learning makes the training of the reachability network unstable when k changes which results in lower performance.

A.2. Effect of tolerance Δt

Adding the tolerance Δt to our k -SP constraint makes it “softer” by allowing Δt -steps of redundancy in transition (See Eq. (18)). Intuitively, a small tolerance may improve the stability of RNet by incorporating a possible noise in RNet prediction, but a very large tolerance will make it less effective in removing sub-optimality in transition. Figure 11 and 12 show the performance of **SPRL** on *DeepMind Lab* and *MiniGrid* domains with varying tolerance Δt . Similar to k , we can see that there exists a “sweet spot” that balances between the reduction in policy space and stabilization of noisy RNet output (*e.g.*, $\Delta t = 25$) in *MiniGrid*. Note that the best tolerance values for *DeepMind Lab* and *MiniGrid* are vastly different. This is mainly because we used *multiple tolerance sampling* (See Appendix E) for *DeepMind Lab* but not for *MiniGrid*. Since the multiple tolerance sampling also improves the stability of RNet, larger tolerance has less benefit compared to its disadvantage.




Figure 11. Average episode reward of **SPRL** with varying $\Delta t = 1, 2, 3, 5$ as a function of environment steps for *DeepMind Lab* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained with $\Delta t = 1$.




Figure 12. Average episode reward of **SPRL** with varying $\Delta t = 10, 15, 25, 50$ as a function of environment steps for *MiniGrid* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained with $\Delta t = 25$.




Figure 13. Average episode reward of **SPRL** with varying observation stacking dimension of 0, 1, 5, 10 as a function of environment steps for *DeepMind Lab* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained without stacking (*i.e.*, #stack=0).




Figure 14. Average episode reward of **SPRL** with varying observation stacking dimension of 0, 1, 5 as a function of environment steps for *MiniGrid* tasks. Other hyper-parameters are kept same as the best hyper-parameter. The best performance is obtained without stacking (*i.e.*, #stack=0).

A.3. Stacking observation

The CMDP with k -SP constraint becomes the $(k + 1)$ -th order MDP as shown in Eq. (13). Thus, in theory, the policy should take current state s_t augmented by stacking the k previous states as input: $[s_{t-k}, s_{t-k+1}, \dots, s_t]$, where $[\cdot]$ is a stacking of the pixel observation along the channel (*i.e.*, color) dimension. However, stacking the observation may not lead to the best empirical results in practice. Figure 13 and 14 show the performance of **SPRL** on *DeepMind Lab* and *MiniGrid* domains with varying stacking dimensions. For stack= m , we stacked the observation from $t - m$ to t : $[s_{t-m}, s_{t-m+1}, \dots, s_t]$. We experimented up to $m = k$: up to $m = 10$ for *DeepMind Lab* and $m = 5$ for *MiniGrid*. The result shows that stacking the

observation does not necessarily improve the performance for MDP order greater than 1, which is often observed when the function approximation is used (e.g., Savinov et al. (2018b)). Thus, we did not augment the observation in all the experiments.




Figure 15. Average episode reward of **SPRL** with varying sparsity controlled by substep $N = 2, 3, 4, 6$ for *Fetch* tasks. Other hyper-parameters are kept same as the best hyper-parameter. As sparsity grows, the difference between **SPRL** and the baselines increase accordingly.

A.4. Effect of Sparsity of the reward

We evaluated **SPRL** on *FetchReach-v1* with a smaller “substep” (i.e., the number of action repetitions) N_{substep} to make the reward sparser, while keeping the episode length \times substep the same to ensure the task is solvable within the episode length. The default N_{substep} is 20, and we used $N_{\text{substep}} = 2, 3, 4, 6$. Figure 15 summarizes the experiment result on varying sparsity of the reward. We can see that as the reward becomes sparser, the performance gap between **SPRL** and the baselines becomes larger, which is consistent with our theory.

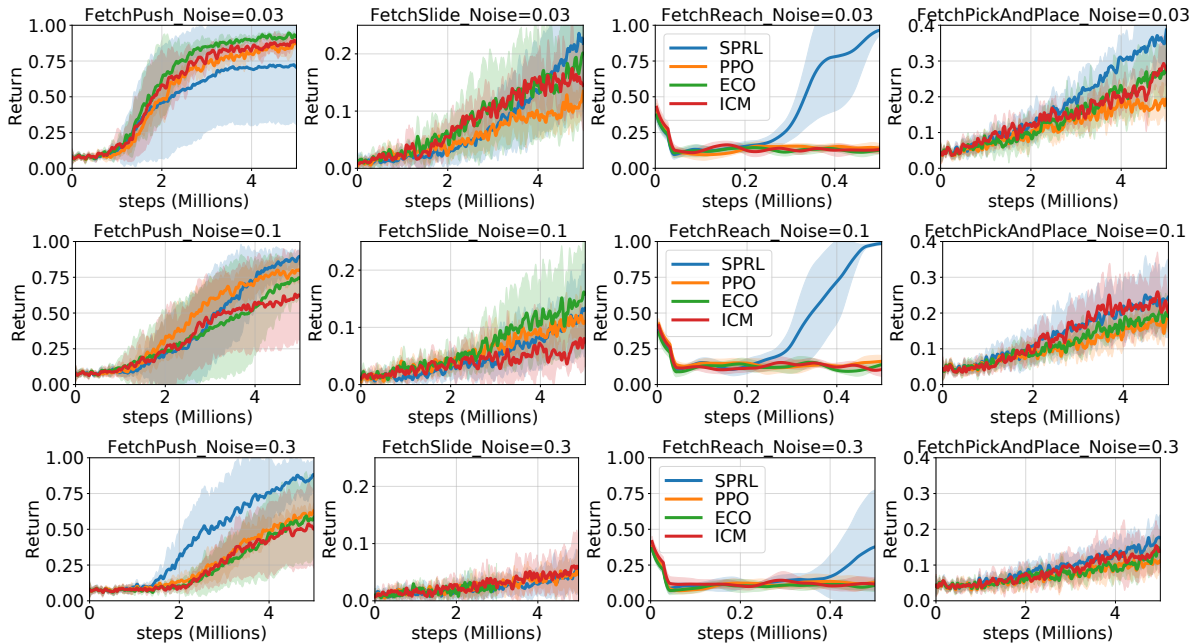


Figure 16. Average episode reward of **SPRL** with varying stochasticity using action noise of $\mathcal{N}(0, \sigma^2)$ where $\sigma = 0.03, 0.1, 0.3$ for *Fetch* tasks. Other hyper-parameters are kept same as the best hyper-parameter. We observe that the performance is preserved the most for **SPRL** when stochasticity changes.

A.5. Effect of stochasticity of the environment

Figure 16 summarizes the experiment result on varying stochasticity with a sparser version of *Fetch* tasks. We used an additive action noise following $\mathcal{N}(0, \sigma^2)$ with $\sigma = 0.03, 0.1, 0.3$ to vary the stochasticity of the environment. Also, we used substep $N_{\text{substep}} = 3, 8, 10, 10$ for *Reach, Push, Slide, Pick* respectively to make the reward sparser. This change was made to clearly observe the performance difference when stochasticity changes. The performance of **SPRL** does not degrade much as the stochasticity increases compared to the other baselines indicating that **SPRL** robustly performs well in stochastic environments.

B. Analysis on the reachability network (RNet)




Figure 17. The accuracy of the learned reachability network on (a) *FourRooms-7×7* (b) *FourRooms-11×11*, (c) *KeyDoors-7×7* and (d) *KeyDoors-11×11* in *MiniGrid* in terms of environment steps.




Figure 18. The accuracy of the learned reachability network on (a) *GoalSmall* (b) *ObjectMany*, and (c) *GoalLarge* in *DeepMind Lab* in terms of environment steps.




Figure 19. The accuracy of the learned reachability network on (a) *Montezuma's Revenge* (b) *Freeway*, (c) *Ms.Pacman*, (d) *Gravitar*, (e) *Seaquest*, and (f) *HERO* in *Atari* in terms of environment steps.

B.1. Accuracy of the reachability network

We measured the accuracy of the reachability network on *MiniGrid*, *DeepMind Lab*, *Atari*, and *Fetch* in Figure 17, Figure 18, Figure 19, and Figure 20. The accuracy was measured on the validation set, where we constructed the validation set by sampling 15,000 positive and negative samples respectively from the replay buffer of size 60,000 for *MiniGrid*, *DeepMind Lab*, and 7,500 positive and negative samples respectively from the replay buffer of size 30,000 for *Atari*, *Fetch*.




Figure 20. The accuracy of the learned reachability network on (a) *FetchPush-v1* (b) *FetchReach-v1*, (c) *FetchSlide-v1* and (d) *FetchPickAndPlace-v1* in *Fetch* in terms of environment steps.

Specifically, for an anchor s_t , we sample the positive sample $s_{t'}$ from $t' \in [t + 1, t + k + \Delta t]$, and the negative sample $s_{t''}$ from $t'' \in [t + k + \Delta t + \Delta_-, t + k + \Delta t + 2\Delta_-]$.

The RNet reaches an accuracy higher than 80% in only 0.4M steps in both *MiniGrid* and *DeepMind Lab*. We note that this is quite high considering the unavoidable noise in the negative samples; since the negative samples are sampled based on the temporal distance, not based on the actual reachability, they have a non-zero probability of being reachable, in which case they are in fact the positive samples. For most of the games of *Atari*, the accuracy of RNet was above 90% except for *Montezuma’s Revenge* where the distribution shift in the state space occurs within a task. We note that RNet maintains a high accuracy above 80% throughout the whole training process due to the additional stabilizing techniques such as weight decay. See Appendix F.3 for more details. Figure 20 summarizes the performance of RNet in *Fetch* tasks. Overall, RNet achieves more than 95% accuracy in less than 1M steps on all four tasks. Considering that predicting the shortest-path distance is much harder in the continuous action domain than the discrete action domains, this result indicates that RNet can be efficiently trained from the time contrastive learning objective even in challenging continuous action domains.

B.2. Ablation study: comparison between the learned RNet and the GT-RNet

In this section, we study the effect of RNet’s accuracy on the SPRL’s performance. To this end, we implement and compare the ground-truth reachability network by computing the ground-truth distance between a pair of states in *MiniGrid*.




Figure 21. The accuracy of the learned reachability network on (a) *FourRooms-7x7* (b) *FourRooms-11x11*, (c) *KeyDoors-7x7* and (d) *KeyDoors-11x11* in *MiniGrid* in terms of environment steps.

Ground-truth reachability network Ground-truth reachability network was implemented by computing the distance between the two-state inputs, and comparing it with k . For the state inputs s and s' , we roll out all possible k -step trajectories starting from the state s using the ground-truth single-step forward model. If s' is ever visited during the roll-out, the output of the k -reachability network is 1 and otherwise, the output is 0.

Result. We compared the performance of our SPRL with the learned RNet and the ground-truth RNet (GT-SPRL) in Figure 21 with the best hyperparameters. Overall, the performance of SPRL and GT-SPRL are similar. This is partly because the learned RNet achieves quite high accuracy in the early stage of learning (see Figure 17). Interestingly, we can observe that our SPRL with learned RNet performs better than SPRL with GT-RNet on *FourRooms-7x7* and *KeyDoors-7x7*. This is possible since a small noise in RNet output can have a similar effect to the increased tolerance Δt on RNet, which makes the resulting cost denser, which may be helpful depending on the tasks and hyperparameters.

C. Qualitative analysis on k -SP cost

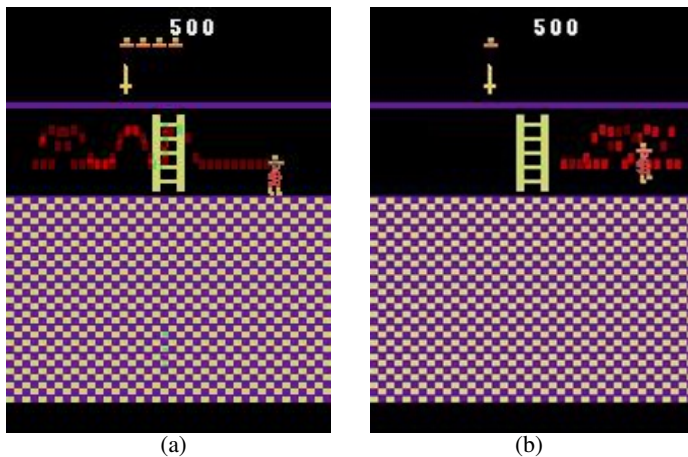


Figure 22. The visualization of k -SP and non- k -SP trajectories in Room 6 of *Montezuma's Revenge*. We visualized the agent's trajectories where it receives (a) low k -SP cost (*i.e.*, lowest 25%) and (b) high k -SP cost (*i.e.*, highest 25%) for 60 steps. The intensity of the red color shows the count of the visited coordinate of the agent. In high k -SP cost trajectory, agent tends to move back and forth between the ladder and the right corner of room 6 while in low k -SP cost trajectory, agent moves forward without redundancy in the transition.

We visually inspected the agent's trajectory with high and low k -SP cost to see whether **SPRL** can correctly differentiate the shortest-path trajectory. Figure 22 (a) is the visualization of the agent's trajectory with low k -SP cost (*i.e.*, shortest-path) trajectory and Figure 22 (b) shows the trajectory with high k -SP cost (*i.e.*, non-shortest path). We can observe that when the agent's trajectory has many redundancies (*e.g.*, moving back-and-forth or staying still in the same place), a high k -SP cost is given to the agent, while the agent is less penalized when it takes the shortest path.

D. Experiment details of *MiniGrid* domain

D.1. Environment

MiniGrid is a 2D grid-world environment with diverse predefined tasks (Chevalier-Boisvert et al., 2018). It has several challenging features such as pictorial observation, random initialization of the agent and the goal, complex action space, and transition dynamics involving the agent’s orientation of movement and changing object status via interaction (e.g., key-door).

State Space. An observation s_t is represented as $H \times W \times C$ tensor, where H and W are the height and width of the map respectively, and C is the features of the objects in the grid. The (h, w) -th element of observation tensor is $(type, color, status)$ of the object and for the coordinate of the agent, the (h, w) -th element is $(type, 0, direction)$. The map size (i.e., $H \times W$) varies depending on the task; e.g., for *FourRooms-7×7* task, the map size is 7×7 .

Action Space and transition dynamics The episode terminates in 100 steps, and the episode may terminate earlier if the agent reaches the goal before 100 steps. The action space consists of seven discrete actions with the following transitions.

- Turn-Counter-Clockwise: change the *direction* counter-clockwise by 90 degree.
- Turn-Clockwise: change the *direction* clockwise by 90 degree.
- Move-Forward: move toward *direction* by 1 step unless blocked by other objects.
- Pick-up-key: pickup the key if the key is in front of the agent.
- Drop-the-key: drop the key in front of the agent.
- Open/Close-doors: open/close the door if the door is in front of the agent.
- Optional-action: not used

Reward function. The reward is given only if the agent reaches the goal location, and the reward magnitude is $1 - 0.9(\text{length of episode}/\text{maximum step for an episode})$. Thus, the agent can maximize the reward by reaching the goal location in the shortest time.

D.2. Tasks

In *FourRooms-7×7* and *FourRooms-11×11*, the map structure has four large rooms, and the agent needs to reach the goal. In *KeyDoors-7×7* and *KeyDoors-11×11*, the agent needs to pick up the key, go to the door, and open the door before reaching the goal location.

D.3. Architecture and hyper-parameters

We used a simple CNN architecture similar to (Mnih et al., 2015) for the policy network. The network consists of Conv1 (16×2×2-1/SAME)-CReLU-Conv2 (8×2×2-1/SAME)-CReLU-Conv3 (8×2×2-1/SAME)-CReLU-FC (512)-FC (action-dimension), where SAME padding ensures the input and output have the same size (i.e., width and height) and CReLU (Shang et al., 2016) is a non-linear activation function applied after each layer. We used Adam (Kingma & Ba, 2014) optimizer to optimize the policy network.

For hyper-parameter search, we swept over a set of hyper-parameters specified in Table 1, and chose the best one in terms of the mean AUC over all the tasks, which is also summarized in Table 1.

E. Experiment details of *DeepMind Lab* domain

E.1. Environment

DeepMind Lab is a 3D-game environment with a first-person view. Along with random initialization of the agent and the goal, complex action space including directional change, random change of texture, color, and maze structure are features that make tasks in *DeepMind Lab* hard to be learned.

PPO		
Hyperparameters	Sweep range	Final value
Learning rate	0.001, 0.002, 0.003	0.003
Entropy	0.003, 0.005, 0.01, 0.02, 0.05	0.01
ICM		
Hyperparameters	Sweep range	Final value
Learning rate	0.001, 0.002, 0.003	0.003
Entropy	-	0.01
Forward/Inverse model loss weight ratio	0.2, 0.5, 0.8, 1.0	0.8
Curiosity module loss weight	0.03, 0.1, 0.3, 1.0	0.3
ICM bonus weight	0.1, 0.3, 1.0, 3.0	0.1
GT-Grid		
Hyperparameters	Sweep range	Final value
Learning rate	0.001, 0.002, 0.003	0.003
Entropy	-	0.01
GT-Grid bonus weight	0.003, 0.01, 0.03, 0.1, 0.3	0.01
ECO		
Hyperparameters	Sweep range	Final value
Learning rate	-	0.003
Entropy	-	0.01
k	3, 5, 7	3
ECO bonus weight	0.001, 0.002, 0.005, 0.01	0.001
SPRL		
Hyperparameters	Sweep range	Final value
Learning rate	0.003, 0.01	0.01
Entropy	-	0.01
k	2, 5	2
Tolerance (Δt)	-	1
Negative bias (Δ^-)	10, 20	20
Positive bias (Δ^+)	-	5
Cost scale (λ)	0.001, 0.002, 0.005	0.002
$N_{\Delta t}$	30, 60	60

Table 1. The range of hyperparameters swept over and the final hyperparameters used in *MiniGrid* domain.

Algorithm 1 Sampling the triplet data from an episode for RNet training**Require:** Hyperparameters: $k \in \mathbb{N}$, Positive bias $\Delta^+ \in \mathbb{N}$, Negative bias $\Delta^- \in \mathbb{N}$

```

1: Initialize  $t_{\text{anc}} \leftarrow 0$ .
2: Initialize  $S_{\text{anc}} = \emptyset, S_+ = \emptyset, S_- = \emptyset$ .
3: while  $t_{\text{anc}} < T$  do
4:    $S_{\text{anc}} = S_{\text{anc}} \cup \{s_{t_{\text{anc}}}\}$ .
5:    $t_+ = \text{Uniform}(t_{\text{anc}} + 1, t_{\text{anc}} + k)$ .
6:    $t_- = \text{Uniform}(t_{\text{anc}} + k + \Delta^-, T)$ .
7:    $S_+ = S_+ \cup \{s_{t_+}\}$ .
8:    $S_- = S_- \cup \{s_{t_-}\}$ .
9:    $t_{\text{anc}} = \text{Uniform}(t_+ + 1, t_+ + \Delta^+)$ .
10: end while
11: Return  $S_{\text{anc}}, S_+, S_-$ 

```

State Space. A state s_t has the dimension of $84 \times 84 \times 3$. The state is given as a first-person view of the map structure. We resized the *RGB* image into $84 \times 84 \times 3$ *RGB* image and normalized it by dividing the pixel value by 255.

Action Space and transition dynamics The episode terminates after the fixed number of steps regardless of the goal being achieved. The original action space consists of seven discrete actions: Move-Forward, Move-Backward, Strafe Left, Strafe Right, Look Left, Look Right, Look Left, and Move-Forward, Look Right and Move-Forward. In our experiment, we used eight discrete actions with the additional action Fire as in Higgins et al. (2017); Vezhnevets et al. (2017); Savinov et al. (2018b); Espeholt et al. (2018); Khetarpal & Precup (2018).

E.2. Tasks

We tested our agent and compared methods on three standard tasks in *DeepMind Lab*: *GoalSmall*, *GoalLarge*, and *ObjectMany* which correspond to `explore_goal_locations_small`, `explore_goal_locations_large`, and `explore_object_rewards_many`, respectively. *GoalSmall* and *GoalLarge* have a single goal in the maze, but the size of the maze is larger in *GoalLarge* than *GoalSmall*. The agent and goal locations are randomly set at the beginning of the episode and the episode length is fixed to 1,350 steps for *GoalSmall* and 1,800 steps for *GoalLarge*. When the agent reaches the goal, it positively rewards the agent and the agent is re-spawned in a random location without terminating the episode, such that the agent can reach the goal multiple times within a single episode. Thus, the agent’s goal is to reach the goal location as many times as possible within the episode length. *ObjectMany* has multiple objects in the maze, where reaching the object positively rewards the agent and the object disappears. The episode length is fixed to 1,800 steps. The agent’s goal is to gather as many objects as possible within the episode length.

E.3. Reachability network Training

Similar to Savinov et al. (2018b), we used the following contrastive loss for training the reachability network:

$$\mathcal{L}_{\text{Rnet}} = -\log(\text{Rnet}_{k-1}(s_{\text{anc}}, s_+)) - \log(1 - \text{Rnet}_{k-1}(s_{\text{anc}}, s_-)), \quad (21)$$

where s_{anc}, s_+, s_- are the anchor, positive, and negative samples, respectively. The anchor, positive and negative samples are sampled from the same episode, and their time steps are sampled according to Algorithm 1. The RNet is trained in an off-policy manner from the replay buffer with the size of 60K environment steps collecting agent’s online experience. We found that adaptive scheduling of RNet is helpful for faster convergence of RNet. Out of 20M total environment steps, for the first 1M, 1M, and 18M environment steps, we updated RNet every 6K, 12K, and 36K environment steps, respectively. For all three environments of *DeepMind Lab*, RNet accuracy was ~ 0.9 after 1M steps.

Multiple tolerance. In order to improve the stability of Reachability prediction, we used the statistics over multiple samples rather than using a single-sample estimate as suggested in Eq. (18). As a choice of sampling method, we simply used multiples of tolerance. In other words, given $s_{t-(k+\Delta t)}$ and s_t as inputs for reachability network, we instead used $s_{t-(k+n\Delta t)}$ and s_t where $1 \leq n \leq N_{\Delta t}$, $n \in \mathbb{N}$ and $N_{\Delta t}$ is the number of tolerance samples. We used 90-percentile of $N_{\Delta t}$ outputs of reachability network, $\text{Rnet}_{k-1}(s_{t-(k+n\Delta t)}, s_t)$, as in (Savinov et al., 2018b) to get the representative of the samples.

Hyperparameters for SPRL	Sweep range	Final value
Learning rate	-	0.0003
Entropy	-	0.004
k	3, 10, 30	10
Tolerance (Δt)	1, 3, 5	1
Negative bias (Δ^-)	5, 10, 20	20
Positive bias (Δ^+)	-	5
Cost scale (λ)	0.02, 0.06, 0.2	0.06
Optimizer	-	Adam
$N_{\Delta t}$	-	200

Table 2. The range of hyperparameters swept over and the final hyperparameters used for our **SPRL** method in *DeepMind Lab* domain.

E.4. Architecture and hyper-parameters

Following (Savinov et al., 2018b), we used the same CNN architecture used in (Mnih et al., 2015).

For **SPRL**, we used a smaller reachability network (RNet) architecture compared to **ECO** to reduce the training time. The RNet is based on the siamese architecture with two branches. Following (Savinov et al., 2018b), **ECO** used Resnet-18 (He et al., 2016) architecture with 2-2-2-2 residual blocks and 512-dimensional output fully connected layer to implement each branch. For **SPRL**, we used Resnet-12 with 2-2-1 residual blocks and 512-dimensional output fully connected layer to implement each branch. The RNet takes two-states as inputs, and each state is fed into each branch. The outputs of the two branches are concatenated and forwarded to three⁴ 512-dimensional fully-connected layers to produce one-dimensional sigmoid output, which predicts the reachability between two state inputs. We also resized the observation to the same dimension as policy (*i.e.*, $84 \times 84 \times 3$, which is smaller than the original $120 \times 160 \times 3$ used in (Savinov et al., 2018b)).

For all the baselines (*i.e.*, **PPO**, **ECO**, **ICM**, and **GT-Grid**), we used the best hyperparameter used in (Savinov et al., 2018b). For **SPRL**, we searched over a set of hyperparameters specified in Table 2, and chose the best one in terms of the mean AUC over all the tasks, which is also summarized in Table 2.

F. Experiment details of *Atari* domain

F.1. Environment

Atari is an important and prominent benchmark in deep reinforcement learning with a high-dimensional visual input. One of the main benefits of using *Atari* as a testbed is that it covers not only navigational tasks but various tasks such as avoiding and destroying enemies by firing a bullet or changing the map structure using bombs as explained in (Bellemare et al., 2013). Because of the diversity of the task *Atari* is covering, solving *Atari* shows that the algorithm has a certain degree of generality. There exists a variety of preprocessing details for *Atari*. We mostly followed the implementation of OpenAI Baselines (Dhariwal et al., 2017). For detailed information, see Table 3.

State Space. A state s_t is represented as $84 \times 84 \times 4$. We stacked 4 consecutive frames achieved by taking the same action 4 times in a row and resized *RGB* image into $84 \times 84 \times 1$ gray image and normalized by dividing the pixel value by 255.

Action Space and transition dynamics The episode terminates when the agent loses all of the lives given. The action space consists of eighteen discrete actions as in (Bellemare et al., 2013).

F.2. Various Tasks of *Atari*: Navigational and Non-navigational

We tested our agent and compared methods on navigational and **non-navigational tasks** in *Atari*: *Montezuma’s Revenge*, *Freeway*, *Ms.Pacman*, *Gravitar*, *Seaquest*, *HERO*. *Montezuma’s Revenge* is a famous game as a hard exploration game in *Atari*. An agent should pick up the items such as a key to open the door or a knife to destroy the enemy. In *Freeway*, an agent should cross the road while avoiding the car. *Ms.Pacman* is a game where an agent should eat the items and avoid the

⁴Savinov et al. (2018b) used four 512-dimensional fully-connected layers.

Parameter	Value
Image Width	84
Image Height	84
Grayscale	Yes
Number of Actions	18
Action Repetitions	4
Frame Stacking	4
End of episode when life lost	No
Reward Clipping	[-1,1]
Discount(γ)	0.99
Max Episode Length	10000
Number of parallel workers	12

Table 3. Preprocessing details for *Atari*

enemy. Three games mentioned until now have a navigational feature meaning that the agent can move toward a certain coordinate to get the score. However, the three games to be mentioned have a non-navigational feature meaning that the agent should not only get to a certain coordinate but also use specific action to get the score. In *Gravitar* and *Seaquest*, an agent should shoot a bullet. In *HERO*, an agent can install a bomb to break the wall and move forward. By adding *Gravitar*, *Seaquest*, and *HERO* to our testbed, we evaluated how **SPRL** performs in non-navigational tasks.

F.3. Reachability network Training

We followed the details of reachability network training for *DeepMind Lab* except for 1) replay buffer size and 2) reachability network training frequency. We changed the size of the replay buffer for the reachability network from 60K environment steps to 30K environment steps. To avoid overfitting of the reachability network, we enlarged the reachability network training frequency from 6K environment steps to 150K environment steps after the initial 1M environment steps.

Stabilizing Reachability Network. In some of the games of *Atari*, within a task exists a distributional shift in the state space that hinders stable reachability network training. Therefore we had to use some techniques to mitigate the instability problem of the reachability network: Weight decay and Label smoothing. We used weight decay with the factor of 0.03 and label smoothing of 0.1. Also instead of using current and future states as inputs, we used current and pixel-level subtraction of current and future states which stabilized the learning.

F.4. Architecture and hyper-parameters

For the policy architecture, we used the same CNN architecture used in Mnih et al. (2015). For the reachability network, we used the same architecture used for *DeepMind Lab* except for the input layer. For the input layer, we used $(s_t, s_t - s_{t-k})$ instead of (s_t, s_{t-k}) . This change helped the reachability network to avoid suffering from overfitting when the distribution shift in the state space occurs.

For hyper-parameter search, we swept over a set of hyper-parameters specified in Table 4 and chose the best one in terms of the mean AUC over all the tasks, which is also summarized in Table 4.

G. Experiment details of *Fetch* domain

G.1. Environment

Fetch is a continuous control environment with a two-fingered gripper. By controlling the gripper, specific interaction between the gripper and the object needs to be accomplished for a given task. The agent and the goal locations are randomly initialized at each episode and these features make the environment difficult to solve.

State Space. At each time step t , the state input s_t is a vector (16-dimensional for *FetchReach-v1* and 31-dimensional for *FetchPush-v1*, *FetchSlide-v1*, and *FetchPickAndPlace-v1*) consisting of the location and the velocity of the gripper. When

PPO		
Hyperparameters	Sweep range	Final value
Learning rate	0.0001, 0.0002, 0.0003, 0.0005	0.0005
Entropy	0.001, 0.003, 0.005, 0.01, 0.03	0.01
ICM		
Hyperparameters	Sweep range	Final value
Learning rate	0.0005	0.0005
Entropy	-	0.01
Forward/Inverse model loss weight ratio	1.0	1.0
Curiosity module loss weight	1.0	1.0
ICM bonus weight	0.0001, 0.0003, 0.001, 0.003, 0.01	0.01
ECO		
Hyperparameters	Sweep range	Final value
Learning rate	0.0001, 0.0003, 0.0005	0.0005
Entropy	-	0.01
ECO bonus weight	0.001, 0.003, 0.01, 0.03, 0.1	0.001
SPRL		
Hyperparameters	Sweep range	Final value
Learning rate	0.0003, 0.0005	0.0005
Entropy	-	0.01
SPRL cost scale (λ)	0.01, 0.03, 0.05, 0.1	0.05
Reachability network (for ECO and SPRL)		
k	5, 8, 12, 15	12
Tolerance (Δt)	-	1
Negative bias (Δ^-)	80, 100, 120	80
Positive bias (Δ^+)	-	5
$N_{\Delta t}$	30, 50, 100, 200, 400	200

Table 4. The range of hyperparameters swept over and the final hyperparameters used in *Atari* domain.

the object exists in the task, the location, rotation, linear and angular velocities of the object are also included in the state s_t .

Action Space The action space is continuous. Actions are total 4 dimensional: 3 dimensions for gripper movement and 1 dimension for the opening of the gripper.

Modification in the environment We made two modifications in the episode termination and the reward function to make it a sparse-reward task. The agent receives +1 reward if the agent reaches the goal and 0 rewards otherwise. Also, the episode terminates when the agent reaches the goal such that the agent can receive a non-zero reward at most once in an episode.

G.2. Tasks

Four tasks are available in the *Fetch* domain: *FetchPush-v1*, *FetchReach-v1*, *FetchSlide-v1*, *FetchPickAndPlace-v1*

- *FetchPush-v1* : The goal is to push a box to a target location. The gripper can only push in this task since fingers are not controllable.
- *FetchReach-v1* : The goal is to move the gripper to a target location. This is the easiest task in *Fetch* since moving the gripper is a fundamental skill required in all four tasks.
- *FetchSlide-v1* : The goal is to hit the object, let the object slide, and stop at the desired location by friction.
- *FetchPickAndPlace-v1* : The goal is to grasp a box and move the box to the target location.

G.3. Reward Normalization

We used reward normalization analogous to [Burda et al. \(2018a\)](#), *i.e.*, “dividing the rewards by a running estimate of the standard deviation of the sum of discounted rewards”. Reward normalization has been particularly effective in the *Fetch* environment on every algorithm. We leave the analysis of this phenomenon as future work.

G.4. Architecture and hyper-parameters

We used a simple MLP architecture for policy network. The network consists of FC (256)-ReLU-FC (256)-ReLU-FC (256)-ReLU-FC (256)-ReLU-FC (action-dimension). The reachability network was also a simple MLP network. The network consists of FC (512)-BatchNorm-ReLU-FC (512)-BatchNorm-ReLU-FC (512)-BatchNorm-ReLU-FC (512)-BatchNorm-ReLU-FC (1)-Sigmoid. We used Adam ([Kingma & Ba, 2014](#)) optimizer to optimize both networks.

For hyper-parameter search, we swept over a set of hyper-parameters specified in [Table 5](#) and chose the best one in terms of the mean AUC over all the tasks, which is also summarized in [Table 5](#).

PPO		
Hyperparameters	Sweep range	Final value
Learning rate	-	0.0005
Entropy	0.001, 0.003, 0.01	0.01
Action Noise	-	0.003
ICM		
Hyperparameters	Sweep range	Final value
Learning rate	-	0.0005
Entropy	-	0.001
Action Noise	-	0.003
Forward/Inverse model loss weight ratio	-	1.0
Curiosity module loss weight	-	1.0
ICM bonus weight	0.0001, 0.0002, 0.0003, 0.0005, 0.001, 0.003, 0.01	0.001
ECO		
Hyperparameters	Sweep range	Final value
Learning rate	-	0.0005
Entropy	-	0.001
Action Noise	-	0.003
ECO bonus weight	0.0001, 0.0002, 0.0003, 0.0005, 0.001, 0.003, 0.01	0.0001
Bias	0.5, 1.0	1.0
SPRL		
Hyperparameters	Sweep range	Final value
Learning rate	-	0.0005
Entropy	-	0.001
SPRL cost scale (λ)	0.0001, 0.0002, 0.0003, 0.0005, 0.001, 0.003, 0.01	0.05
Bias	0.5, 1.0	1.0
Reachability network (for ECO and SPRL)		
k	-	10
Tolerance (Δt)	-	1
Negative bias (Δ^-)	1, 2, 3, 5, 8, 12	12
Positive bias (Δ^+)	-	5
$N_{\Delta t}$	-	20

Table 5. The range of hyperparameters swept over and the final hyperparameters used in *Fetch* domain.

H. Option framework-based formulation

H.1. Preliminary: option framework

Options framework (Sutton, 1998) defines options as a generalization of actions to include a temporally extended series of action. Formally, options consist of three components: a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, a termination condition $\beta : \mathcal{S}^+ \rightarrow [0, 1]$, and an initiation set $\mathcal{I} \subseteq \mathcal{S}$. An option $\langle \mathcal{I}, \pi, \beta \rangle$ is available in state s if and only if $s \in \mathcal{I}$. If the option is taken, then actions are selected according to π until the option terminates stochastically according to β . Then, the option-reward and option-transition models are defined as

$$r_s^o = \mathbb{E} \{ r_{t+1} + \gamma r_{t+2} + \dots + \gamma^{k-1} r_{t+k} \mid E(o, s, t) \} \quad (22)$$

$$P_{ss'}^o = \sum_{k=1}^{\infty} p(s', k) \gamma^k \quad (23)$$

where $t+k$ is the random time at which option o terminates, $E(o, s, t)$ is the event that option o is initiated in state s at time t , and $p(s', k)$ is the probability that the option terminates in s' after k steps. Using the option models, we can re-write Bellman equation as follows:

$$V^\pi(s) = \mathbb{E} [r_{t+1} + \dots + \gamma^{k-1} r_{t+k} + \gamma^k V^\pi(s_{t+k})], \quad (24)$$

$$= \sum_{o \in \mathcal{O}} Pr[E(o, s)] \left[r_s^o + \sum_{s'} P_{ss'}^o V^\pi(s') \right]. \quad (25)$$

where $t+k$ is the random time at which option o terminates and $E(o, s)$ is the event that option o is initiated in state s .

H.2. Option-based view-point of shortest-path constraint

In this section, we present an option framework-based viewpoint of our shortest-path (SP) constraint. We will first show that a (sparse-reward) MDP can be represented as a weighted directed graph where nodes are rewarding states, and edges are options. Then, we show that a policy satisfying SP constraint also maximizes the option-transition probability $P_{ss'}^o$.

For a given MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \rho, \bar{\mathcal{S}})$, let $\mathcal{S}^R = \{s \mid R(s) \neq 0\} \subset \mathcal{S}$ be the set of all rewarding states, where $R(s)$ is the reward function upon arrival to state s . In sparse-reward tasks, it is assumed that $|\mathcal{S}^R| \ll |\mathcal{S}|$. Then, we can form a weighted directed graph $G^\pi = (\mathcal{V}, \mathcal{E})$ of policy π and given MDP. The vertex set is defined as $\mathcal{V} = \mathcal{S}^R \cup \rho_0 \cup \bar{\mathcal{S}}$ where \mathcal{S}^R is rewarding states, ρ_0 is the initial states, and $\bar{\mathcal{S}}$ is the terminal states. Similar to the path set in Definition 2, let $\mathcal{T}_{s \rightarrow s'}$ denotes a set of paths transitioning from one vertex $s \in \mathcal{V}$ to another vertex $s' \in \mathcal{V}$:

$$\mathcal{T}_{s \rightarrow s'} = \{ \tau \mid s_0 = s, s_{\ell(\tau)} = s', \{s_t\}_{0 < t < \ell(\tau)} \cap \mathcal{V} = \emptyset \}. \quad (26)$$

Then, the edge from a vertex $s \in \mathcal{V}$ to another vertex $s' \in \mathcal{V}$ is defined by an (implicit) option tuple: $o(s, s') = (\mathcal{I}, \pi, \beta)_{(s, s')}$, where $\mathcal{I} = \{s\}$, $\beta(s) = \mathbb{I}(s = s')$, and

$$\pi^{(s, s')}(\tau) = \begin{cases} \frac{1}{Z} \pi(\tau) & \text{for } \tau \in \mathcal{T}_{s \rightarrow s'} \\ 0 & \text{otherwise} \end{cases}, \quad (27)$$

where Z is the partition function to ensure $\int \pi^{(s, s')}(\tau) d\tau = 1$. Following Eq. (22), the option-reward is given as

$$r_{s, s'}^\pi = \mathbb{E}^{\pi^{(s, s')}} \left[r_{t+1} + \gamma r_{t+2} + \dots + \gamma^{k-1} r_{t+k} \mid E(o^{(s, s')}, s, t) \right], \quad (28)$$

$$= \mathbb{E}^{\pi^{(s, s')}} \left[\gamma^{k-1} r_{t+k} \mid E(o^{(s, s')}, s, t) \right], \quad (29)$$

where $t+k$ is the random time at which option $o(s, s')$ terminates, and $E(o, s, t)$ is the event that option $o(s, s')$ is initiated in state s at time t . Note that in the last equality, $r_{t+1} = \dots = r_{t+k-1} = 0$ holds since $\{s_{t+1}, \dots, s_{t+k-1}\} \cap \mathcal{V} = \emptyset$ from

the definition of option policy $\pi^{(s,s')}$. Following Eq. (23), the option transition is given as

$$P_{s,s'}^\pi = \sum_{k=1}^{\infty} p(s', k) \gamma^k \quad (30)$$

$$= \mathbb{E}^\pi [\gamma^k | s_0 = s, s_k = s', \{r_t\}_{t < k} = 0] \quad (31)$$

$$= \gamma^{D_{\text{nr}}^\pi(s,s')}. \quad (32)$$

where $p(s', k)$ is the probability that the option terminates in s' after k steps, and $D_{\text{nr}}^\pi(s, s')$ is the π -distance in Definition 3. Then, we can re-write the shortest-path constraint in terms of $P_{s,s'}^\pi$ as follows:

$$\Pi^{\text{SP}} = \{\pi | \forall (s, s' \in \mathcal{T}_{\hat{s}, \hat{s}', \text{nr}}^\pi \text{ s.t. } (\hat{s}, \hat{s}') \in \Phi^\pi), D_{\text{nr}}^\pi(s, s') = \min_{\pi} D_{\text{nr}}^\pi(s, s')\} \quad (33)$$

$$= \{\pi | \forall (s, s' \in \mathcal{T}_{\hat{s}, \hat{s}', \text{nr}}^\pi \text{ s.t. } (\hat{s}, \hat{s}') \in \Phi^\pi), P_{s,s'}^\pi = \max_{\pi} P_{s,s'}^\pi\} \quad (34)$$

Thus, we can see that the policy satisfying SP constraint also maximizes the option-transition probability. We will use this result in Appendix I.

I. Shortest-Path Constraint: A Single-goal Case

In this section, we provide more discussion on a special case of the shortest-path constraint (Section 3.1), when the (stochastic) MDP defines a single-goal task: *i.e.*, there exists a unique initial state $s_{\text{init}} \in \mathcal{S}$ and a unique goal state $s_g \in \mathcal{S}$ such that s_g is a terminal state, and $R(s) > 0$ if and only if $s = s_g$.

We first note that the non-rewarding path set is identical to the path set in such a setting, because the condition $r_t = 0 (t < \ell(\tau))$ from Definition 2 is always satisfied as $R(s) > 0 \Leftrightarrow s = s_g$ and $s_{\ell(\tau)} = s_g$:

$$\mathcal{T}_{s,s',\text{nr}}^\pi = \mathcal{T}_{s,s'}^\pi = \{\tau | s_0 = s, s_{\ell(\tau)} = s', p_\pi(\tau) > 0, \{s_t\}_{t < \ell(\tau)} \neq s'\} \quad (35)$$

Again, $\mathcal{T}_{s,s'}^\pi$ is a set of all path starting from s (*i.e.*, and ending at s' (*i.e.*, $s_{\ell(\tau)} = s'$) where the agent visits s' *only* at the end (*i.e.*, $\{s_t\}_{t < \ell(\tau)} \neq s'$), that can be rolled out by policy with a non-zero probability (*i.e.*, $p_\pi(\tau) > 0$).

We now claim that an optimal policy satisfies the shortest-path constraint. The idea is that, since s_g is the only rewarding and terminal state, maximizing $R(\tau) = \gamma^T R(s_g)$ where $s_T = s_g$ corresponds to minimizing the number of time steps T to reach s_g . In this setting, a shortest-path policy is indeed optimal.

Lemma 4. *For a single-goal MDP, any optimal policy satisfies the shortest-path constraint.*

Proof. Let s_{init} be the initial state and s_g be the goal state. We will prove that any optimal policy is a shortest-path policy from the initial state to the goal state. We use the fact that s_g is the only rewarding state, *i.e.*, $R(s) > 0$ entails $s = s_g$.

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{s \sim \rho}^{\tau \sim \pi} \left[\sum_t \gamma^t r_t \mid s_0 = s \right] \quad (36)$$

$$= \arg \max_{\pi} \mathbb{E}^{\tau \sim \pi} \left[\sum_t \gamma^t r_t \mid s_0 = s_{\text{init}} \right] \quad (37)$$

$$= \arg \max_{\pi} \mathbb{E}^{\tau \sim \pi} [\gamma^T R(s_g) \mid s_0 = s_{\text{init}}, s_{\ell(\tau)} = s_g] \quad (38)$$

$$= \arg \max_{\pi} \mathbb{E}^{\tau \sim \pi} [\gamma^T \mid s_0 = s_{\text{init}}, s_{\ell(\tau)} = s_g] \quad (39)$$

$$= \arg \min_{\pi} \log_{\gamma} (\mathbb{E}^{\tau \sim \pi} [\gamma^T \mid s_0 = s_{\text{init}}, s_{\ell(\tau)} = s_g]) \quad (40)$$

$$= \arg \min_{\pi} D_{\text{nr}}^\pi(s_{\text{init}}, s_g), \quad (41)$$

where Eq. (39) holds since $R(s_g) > 0$ from our assumption that $R(s) + V^*(s) > 0$. \square

J. Proof of Theorem 1

We make the following assumptions on the Markov Decision Process (MDP) \mathcal{M} : namely *mild stochasticity* (Definitions 8 and 9).

Definition 8 (Mild stochasticity (1)). *In MDP \mathcal{M} , there exists an optimal policy π^* and the corresponding shortest-path policy $\pi^{sp} \in \Pi^{SP}$ such that for all $s, s' \in \Phi^\pi$, it holds $p_{\pi^*}(\bar{s} = s' | s_0 = s) = p_{\pi^{sp}}(\bar{s} = s' | s_0 = s)$.*

Definition 9 (Mild stochasticity (2)). *In MDP \mathcal{M} , the optimal policy π^* does not visit the same state more than once: For all $s \in \mathcal{S}$ such that $\rho_{\pi^*}(s) > 0$, it holds $\rho_{\pi^*}(s) = 1$, where $\rho_\pi(s) \triangleq \mathbb{E}_{s_0 \sim \rho_0(S), a \sim \pi(A|s), s' \sim (S|s,a)} \left[\sum_{t=1}^T \mathbb{I}(s_t = s) \right]$ is the state-visitation count.*

In other words, we assume that the optimal policy does not have a cycle. One common property of MDP that meets this condition is that the reward disappearing after being acquired by the agent. We note that this assumption holds for many practical environments. In fact, in many cases as well as *Atari*, *DeepMind Lab*, etc.

Theorem 1. *For any MDP with the mild stochasticity condition, an optimal policy π^* satisfies the shortest-path constraint: $\pi^* \in \Pi^{SP}$.*

Proof. For simplicity, we prove this based on the option-based view point (see Appendix H). By plugging Eq. (29) and Eq. (31) into Eq. (25), we can re-write the Bellman equation of the value function $V^\pi(s)$ as follows:

$$V^\pi(s) = \sum_{o \in \mathcal{O}} Pr[E(o, s)] \left[r_s^o + \sum_{s'} P_{ss'}^o V^\pi(s') \right] \quad (42)$$

$$= \sum_{s' \in \mathcal{S}^{IR}} p_\pi(\bar{s} = s' | s_0 = s) \left[R(s') \mathbb{E}^{\tau \sim \pi}(\gamma^{\ell(\tau)} | s_0 = s, \bar{s} = s') + \gamma P_{s,s'}^\pi V^\pi(s') \right] \quad (43)$$

$$= \sum_{s' \in \mathcal{S}^{IR}} p_\pi(\bar{s} = s' | s_0 = s) \left[R(s') P_{s,s'}^\pi + \gamma P_{s,s'}^\pi V^\pi(s') \right], \quad (44)$$

$$= \sum_{s' \in \mathcal{S}^{IR}} p_\pi(\bar{s} = s' | s_0 = s) P_{s,s'}^\pi [R(s') + \gamma V^\pi(s')], \quad (45)$$

where \bar{s} is the first rewarding state that agent encounters. Intuitively, $p_\pi(\bar{s} = s' | s_0 = s)$ means the probability that the s' is the first rewarding state that policy π encounters when it starts from s . From Eq. (34), our goal is to show:

$$\pi^* \in \Pi^{SP} = \{ \pi \mid \forall (s, s') \in \mathcal{T}_{\Phi, nr}^\pi, P_{s,s'}^\pi = P_{s,s'}^* \}, \quad (46)$$

where $P_{s,s'}^* = \max_\pi P_{s,s'}^\pi$.

We will prove Eq. (46) by contradiction. Suppose π^* is an optimal policy such that $\pi^* \notin \Pi^{SP}$. Then,

$$\exists (\hat{s}, \hat{s}' \in \mathcal{T}_{\Phi, nr}^{\pi^*}) \text{ s.t. } P_{\hat{s}, \hat{s}'}^{\pi^*} \neq P_{\hat{s}, \hat{s}'}^*. \quad (47)$$

Recall the definition: $P_{s,s'}^* = \max_\pi P_{s,s'}^\pi$. Then, for any π , the following statement is true.

$$P_{s,s'}^\pi \neq P_{s,s'}^* \Leftrightarrow P_{s,s'}^\pi < P_{s,s'}^*. \quad (48)$$

Thus, we have

$$P_{\hat{s}, \hat{s}'}^{\pi^*} < P_{\hat{s}, \hat{s}'}^* \quad (49)$$

Let $\pi_{sp} \in \Pi^{SP}$ be a shortest path policy that preserves stochastic dynamics from Definition 8. Then, we have

$$P_{\hat{s}, \hat{s}'}^{\pi_{sp}} < P_{\hat{s}, \hat{s}'}^* = P_{\hat{s}, \hat{s}'}^{\pi_{sp}}. \quad (50)$$

Then, let's compose a new policy $\hat{\pi}$:

$$\hat{\pi}(a|s) = \begin{cases} \pi_{sp}(a|s) & \text{if } \exists \tau \in \mathcal{T}_{\hat{s}, \hat{s}', nr}^{\pi_{sp}} \text{ s.t. } s \in \tau \\ \pi^*(a|s) & \text{otherwise} \end{cases}. \quad (51)$$

Now consider a path $\tau_{\hat{s} \rightarrow \hat{s}'}$ that agent visits \hat{s} at time $t = i$ and transitions to \hat{s}' at time $t = j > i$ while not visiting any rewarding state from $t = i$ to $t = j$ with non-zero probability (i.e., $p_{\pi_{sp}}(\tau) > 0$). We can define a set of such paths as follows:

$$\hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'} = \{ \tau \mid \exists (i < j), s_i = \hat{s}, s_j = \hat{s}', \{s_t\}_{i < t < j} \cap \mathcal{S}^{IR} = \emptyset, p_{\pi_{sp}}(\tau) > 0 \}. \quad (52)$$

To reiterate the definitions from Definition 6: $\mathcal{S}^{\text{IR}} = \{s \mid R(s) > 0 \text{ or } \rho(s) > 0\}$ is the union of all initial and rewarding states, and $\Phi^\pi = \{(s, s') \mid s, s' \in \mathcal{S}^{\text{IR}}, \rho(s) > 0, \mathcal{T}_{s, s', \text{nr}}^\pi \neq \emptyset\}$ is the subset of \mathcal{S}^{IR} such that agent may roll out.

From Definition 9 and Eq. (51), the likelihood of a path τ under policy $\hat{\pi}$ is given as follows:

$$p_{\hat{\pi}}(\tau) = \begin{cases} p_{\pi^*}(\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) p_{\pi_{\text{sp}}}(\tau \mid \tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) & \text{for } \tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'} \\ p_{\pi^*}(\tau) & \text{otherwise} \end{cases}, \quad (53)$$

where $p_{\hat{\pi}}(\tau)$ is the likelihood of trajectory τ under policy $\hat{\pi}$, $p_{\hat{\pi}}(\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) = \int_{\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'} } p_{\hat{\pi}}(\tau) d\tau$ ensures the likelihood $\hat{\pi}(\tau)$ to be a valid probability density function (i.e., $\int p_{\hat{\pi}}(\tau) d\tau = 1$). From the path $\tau_{\hat{s} \rightarrow \hat{s}'}$ and i, j , we will choose two states $s_{\text{ir}}, s'_{\text{ir}} \sim \tau_{\hat{s} \rightarrow \hat{s}'}$, where

$$s_{\text{ir}} = \max_t(s_t \mid s_t \in \mathcal{S}^{\text{IR}}, t \leq i), \quad s'_{\text{ir}} = \min_t(s_t \mid s_t \in \mathcal{S}^{\text{IR}}, j \leq t). \quad (54)$$

Note that such s_{ir} and s'_{ir} always exist in $\tau_{\hat{s} \rightarrow \hat{s}'}$ since the initial state and the terminal state satisfy the condition to be s_{ir} and s'_{ir} .

Then, we can show that the path between s_{ir} and s'_{ir} is **not** a shortest-path. Recall the definition of $D_{\text{nr}}^\pi(s, s')$ (Definition 3):

$$D_{\text{nr}}^{\pi^*}(s_{\text{ir}}, s'_{\text{ir}}) := \log_\gamma \left(\mathbb{E}_{\tau \sim \pi^*} \left[\gamma^{\ell(\tau)} \mid \tau \in \mathcal{T}_{s_{\text{ir}}, s'_{\text{ir}}, \text{nr}}^{\pi^*} \right] \right) \quad (55)$$

$$= \log_\gamma \left(\mathbb{E}_{\tau \sim \pi^*} \left[\underbrace{\gamma^{\ell(\tau)} \mid \tau \in \mathcal{T}_{s_{\text{ir}}, s'_{\text{ir}}, \text{nr}}^{\pi^*}}_{\clubsuit} \right] \right) \quad (56)$$

where we will use $\clubsuit := \gamma^{\ell(\tau)} \mid \tau \in \mathcal{T}_{s_{\text{ir}}, s'_{\text{ir}}, \text{nr}}^{\pi^*}$ for a shorthand notation. Then, we have

$$\gamma^{D_{\text{nr}}^{\pi^*}(s_{\text{ir}}, s'_{\text{ir}})} := \mathbb{E}_{\tau \sim \pi^*} [\clubsuit] \quad (57)$$

$$= p_{\pi^*}(\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \pi^*} [\clubsuit \mid \tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \\ + p_{\pi^*}(\tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \pi^*} [\clubsuit \mid \tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \quad (58)$$

$$\text{(From Definition 5)} < p_{\pi^*}(\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \pi_{\text{sp}}} [\clubsuit \mid \tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \\ + p_{\pi^*}(\tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \pi^*} [\clubsuit \mid \tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \quad (59)$$

$$\text{(From Eq. (53))} = p_{\hat{\pi}}(\tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \hat{\pi}} [\clubsuit \mid \tau \in \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \\ + p_{\hat{\pi}}(\tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}) \mathbb{E}_{\tau \sim \hat{\pi}} [\clubsuit \mid \tau \notin \hat{\mathcal{T}}_{\hat{s} \rightarrow \hat{s}'}] \quad (60)$$

$$= \mathbb{E}_{\tau \sim \hat{\pi}} [\clubsuit] = \gamma^{D_{\text{nr}}^{\hat{\pi}}(s_{\text{ir}}, s'_{\text{ir}})} \quad (61)$$

$$\iff D_{\text{nr}}^{\pi^*}(s_{\text{ir}}, s'_{\text{ir}}) > D_{\text{nr}}^{\hat{\pi}}(s_{\text{ir}}, s'_{\text{ir}}) \quad (62)$$

where Ineq. (62) is given by the fact that $\gamma < 1$. Then, $P_{s_{\text{ir}}, s'_{\text{ir}}}^{\pi^*} < P_{s_{\text{ir}}, s'_{\text{ir}}}^{\hat{\pi}}$.

From Eq. (45), we have

$$V^{\hat{\pi}}(s_{\text{ir}}) = \sum_{s' \in \mathcal{S}^{\text{IR}}} p_{\hat{\pi}}(\bar{s} = s' \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'}^{\hat{\pi}} [R(s') + \gamma V^{\hat{\pi}}(s')] \quad (63)$$

$$\begin{aligned} &= p_{\hat{\pi}}(\bar{s} = s'_{\text{ir}} \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'_{\text{ir}}}^{\hat{\pi}} [R(s'_{\text{ir}}) + \gamma V^{\hat{\pi}}(s'_{\text{ir}})] \\ &\quad + \sum_{s' \in \mathcal{S}^{\text{IR}} \setminus s'_{\text{ir}}} p_{\hat{\pi}}(\bar{s} = s' \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'}^{\hat{\pi}} [R(s') + \gamma V^{\hat{\pi}}(s')] \end{aligned} \quad (64)$$

$$\begin{aligned} &= p_{\pi^*}(\bar{s} = s'_{\text{ir}} \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'_{\text{ir}}}^{\hat{\pi}} [R(s'_{\text{ir}}) + \gamma V^{\pi^*}(s'_{\text{ir}})] \\ &\quad + \sum_{s' \in \mathcal{S}^{\text{IR}} \setminus s'_{\text{ir}}} p_{\pi^*}(\bar{s} = s' \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'}^{\pi^*} [R(s') + \gamma V^{\pi^*}(s')] \end{aligned} \quad (65)$$

$$\begin{aligned} &> p_{\pi^*}(\bar{s} = s'_{\text{ir}} \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'_{\text{ir}}}^{\pi^*} [R(s'_{\text{ir}}) + \gamma V^{\pi^*}(s'_{\text{ir}})] \\ &\quad + \sum_{s' \in \mathcal{S}^{\text{IR}} \setminus s'_{\text{ir}}} p_{\pi^*}(\bar{s} = s' \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'}^{\pi^*} [R(s') + \gamma V^{\pi^*}(s')] \end{aligned} \quad (66)$$

$$= \sum_{s' \in \mathcal{S}^{\text{IR}}} p_{\pi^*}(\bar{s} = s' \mid s_0 = s_{\text{ir}}) P_{s_{\text{ir}}, s'}^{\pi^*} [R(s') + \gamma V^{\pi^*}(s')] \quad (67)$$

$$= V^*(s_{\text{ir}}), \quad (68)$$

where Eq. (65) holds from the *mild-stochasticity (1)* and *mild-stochasticity (2)* assumption, and Ineq. (66) holds because $P_{s_{\text{ir}}, s'_{\text{ir}}}^{\hat{\pi}} > P_{s_{\text{ir}}, s'_{\text{ir}}}^{\pi^*}$ and $R(s') + \gamma V^{\pi^*}(s') > 0$ from the non-negative optimal value assumption (See Section 2). Finally, this is a contradiction since the optimal value function $V^*(s)$ should be the maximum. \square

K. Extended related works

Approximate state abstraction. The approximate state abstraction approaches investigate partitioning an MDP’s state space into clusters of similar states while preserving the optimal solution. Researchers have proposed several state similarity metrics for MDPs. Dean et al. (2013) proposed to use the bisimulation metrics (Givan et al., 2003; Ferns et al., 2004), which measures the difference in transition and reward function. Bertsekas et al. (1988) used the magnitude of Bellman residual as a metric. Abel et al. (2016; 2018); Li et al. (2006) used the different types of distance in optimal Q-value to measure the similarity between states to bound the sub-optimality in optimal value after the abstraction. Recently, Castro (2019) extended the bisimulation metrics to the approximate version for the deep-RL setting where the tabular representation of state is not available.

Our shortest-path constraint can be seen as a form of state abstraction, in that, ours also aim to reduce the size of MDP (*i.e.*, state and action space) while preserving the “solution quality”. However, our method does so by removing sub-optimal policies, not by aggregating similar states (or policies).

Connection to Option framework Our shortest-path constraint constrains the policy space to a set of shortest-path policies (See Definition 5 for definition) between initial and rewarding states. It can be seen as a set of options (Sutton, 1998) transitioning between initial and rewarding states. We refer the readers to Appendix H for the detailed description of the option framework-based formulation of our framework.

References

- Abel, D., Hershkowitz, D., and Littman, M. Near optimal behavior via approximate state abstraction. In *International Conference on Machine Learning*, pp. 2915–2923, 2016.
- Abel, D., Arumugam, D., Lehnert, L., and Littman, M. State abstractions for lifelong reinforcement learning. In *International Conference on Machine Learning*, pp. 10–19, 2018.
- Altman, E. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

- Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 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.
- Bellemare, M. G., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. *arXiv preprint arXiv:1606.01868*, 2016.
- Bellman, R. On a routing problem. *Quarterly of applied mathematics*, 16(1):87–90, 1958.
- Bertsekas, D. P. and Tsitsiklis, J. N. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- Bertsekas, D. P. and Tsitsiklis, J. N. Neuro-dynamic programming: an overview. In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, volume 1, pp. 560–564. IEEE, 1995.
- Bertsekas, D. P., Castanon, D. A., et al. Adaptive aggregation methods for infinite horizon dynamic programming. 1988.
- Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., and Efros, A. A. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018a.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018b.
- Castro, P. S. Scalable methods for computing state similarity in deterministic markov decision processes. *arXiv preprint arXiv:1911.09291*, 2019.
- Chevalier-Boisvert, M., Willems, L., and Pal, S. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Choi, J., Guo, Y., Moczulski, M., Oh, J., Wu, N., Norouzi, M., and Lee, H. Contingency-aware exploration in reinforcement learning. *arXiv preprint arXiv:1811.01483*, 2018.
- Coulom, R. Efficient selectivity and backup operators in monte-carlo tree search. In *International conference on computers and games*, pp. 72–83. Springer, 2006.
- Dean, T. L., Givan, R., and Leach, S. Model reduction techniques for computing approximately optimal solutions for markov decision processes. *arXiv preprint arXiv:1302.1533*, 2013.
- Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y., and Zhokhov, P. Openai baselines. <https://github.com/openai/baselines>, 2017.
- 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.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- Ferns, N., Panangaden, P., and Precup, D. Metrics for finite markov decision processes. In *UAI*, volume 4, pp. 162–169, 2004.
- Florensa, C., Held, D., Wulfmeier, M., Zhang, M., and Abbeel, P. Reverse curriculum generation for reinforcement learning. In *Conference on robot learning*, pp. 482–495. PMLR, 2017.
- Ford Jr, L. R. Network flow theory. Technical report, Rand Corp Santa Monica Ca, 1956.
- Givan, R., Dean, T., and Greig, M. Equivalence notions and model minimization in markov decision processes. *Artificial Intelligence*, 147(1-2):163–223, 2003.

- Guo, X., Singh, S., Lewis, R., and Lee, H. Deep learning for reward design to improve monte carlo tree search in atari games. *arXiv preprint arXiv:1604.07095*, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Higgins, I., Pal, A., Rusu, A., Matthey, L., Burgess, C., Pritzel, A., Botvinick, M., Blundell, C., and Lerchner, A. Darla: Improving zero-shot transfer in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1480–1490. JMLR. org, 2017.
- Huang, Z., Liu, F., and Su, H. Mapping state space using landmarks for universal goal reaching. In *Advances in Neural Information Processing Systems*, pp. 1940–1950, 2019.
- Kakade, S. M. et al. *On the sample complexity of reinforcement learning*. PhD thesis, University of London London, England, 2003.
- Khetarpal, K. and Precup, D. Attend before you act: Leveraging human visual attention for continual learning. *arXiv preprint arXiv:1807.09664*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lai, T. L. and Robbins, H. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Laskin, M., Emmons, S., Jain, A., Kurutach, T., Abbeel, P., and Pathak, D. Sparse graphical memory for robust planning. *arXiv preprint arXiv:2003.06417*, 2020.
- Li, L., Walsh, T. J., and Littman, M. L. Towards a unified theory of state abstraction for mdps. In *ISAIM*, 2006.
- 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.
- Nachum, O., Gu, S. S., Lee, H., and Levine, S. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 3303–3313, 2018.
- Ng, A. Y., Harada, D., and Russell, S. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pp. 278–287, 1999.
- Norvig, P. R. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2002.
- Oh, J., Guo, Y., Singh, S., and Lee, H. Self-imitation learning. In *International Conference on Machine Learning*, pp. 3878–3887. PMLR, 2018.
- Oudeyer, P.-Y. and Kaplan, F. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2778–2787. JMLR. org, 2017.
- Plappert, M., Andrychowicz, M., Ray, A., McGrew, B., Baker, B., Powell, G., Schneider, J., Tobin, J., Chociej, M., Welinder, P., et al. Multi-goal reinforcement learning: Challenging robotics environments and request for research. *arXiv preprint arXiv:1802.09464*, 2018.
- Riedmiller, M., Hafner, R., Lampe, T., Neunert, M., Degraeve, J., Wiele, T., Mnih, V., Heess, N., and Springenberg, J. T. Learning by playing solving sparse reward tasks from scratch. In *International Conference on Machine Learning*, pp. 4344–4353. PMLR, 2018.
- Savinov, N., Dosovitskiy, A., and Koltun, V. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*, 2018a.
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T., and Gelly, S. Episodic curiosity through reachability. *ICLR*, 2018b.

- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal value function approximators. In *International conference on machine learning*, pp. 1312–1320, 2015.
- Schmidhuber, J. Adaptive confidence and adaptive curiosity. In *Institut fur Informatik, Technische Universitat Munchen, Arcisstr. 21, 800 Munchen 2*. Citeseer, 1991.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shang, W., Sohn, K., Almeida, D., and Lee, H. Understanding and improving convolutional neural networks via concatenated rectified linear units. In *international conference on machine learning*, pp. 2217–2225, 2016.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Sutton, R. S. Between mdps and semi-mdps: Learning, planning, and representing knowledge at multiple temporal scales. 1998.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Tessler, C., Mankowitz, D. J., and Mannor, S. Reward constrained policy optimization. *ICLR*, 2019.
- Vezhnevets, A. S., Osindero, S., Schaul, T., Heess, N., Jaderberg, M., Silver, D., and Kavukcuoglu, K. Feudal networks for hierarchical reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 3540–3549. JMLR. org, 2017.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Vodopivec, T., Samothrakis, S., and Ster, B. On monte carlo tree search and reinforcement learning. *Journal of Artificial Intelligence Research*, 60:881–936, 2017.
- Zhang, A., Lerer, A., Sukhbaatar, S., Fergus, R., and Szlam, A. Composable planning with attributes. *ICML*, 2018.
- Zhang, T., Guo, S., Tan, T., Hu, X., and Chen, F. Generating adjacency-constrained subgoals in hierarchical reinforcement learning. *arXiv preprint arXiv:2006.11485*, 2020.