

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

Public Domain Mark 1.0

<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Towards an Interpretable Hierarchical Agent Framework using Semantic Goals

Bharat Prakash¹, Nicholas Waytowich², Tim Oates¹, Tinoosh Mohsenin¹,

¹University of Maryland, Baltimore County, ²US Army Research Lab

Abstract

Learning to solve long horizon temporally extended tasks with reinforcement learning has been a challenge for several years now. We believe that it is important to leverage both the hierarchical structure of complex tasks and to use expert supervision whenever possible to solve such tasks. This work introduces an interpretable hierarchical agent framework by combining planning and semantic goal directed reinforcement learning. We assume access to certain spatial and haptic predicates and construct a simple and powerful semantic goal space. These semantic goal representations are more interpretable, making expert supervision and intervention easier. They also eliminate the need to write complex, dense reward functions thereby reducing human engineering effort. We evaluate our framework on a robotic block manipulation task and show that it performs better than other methods, including both sparse and dense reward functions. We also suggest some next steps and discuss how this framework makes interaction and collaboration with humans easier.

Introduction

Deep reinforcement learning has been successful in many tasks, including robotic control, games, energy management, etc. (Mnih et al. 2015; Schulman et al. 2017; Warnell et al. 2018). However, it has many challenges, such as exploration under sparse rewards, generalization, safety, etc. This makes it difficult to learn good policies in a sample efficient way. Popular ways to tackle these problems include using expert feedback (Christiano et al. 2017; Warnell et al. 2018; Prakash et al. 2020) and leveraging the hierarchical structure of complex tasks. There is a long list of prior work which learns hierarchical policies to break down tasks into smaller sub-tasks (Sutton, Precup, and Singh 1999; Fruit and Lazaric 2017; Bacon, Harb, and Precup 2017; Prakash et al. 2021). Some of them discover options or sub-tasks in an unsupervised way. On the other hand, using some form of supervision, either by providing details about the sub-tasks, intermediate rewards or high-level guidance is a recent approach (Prakash et al. 2021) (Jiang et al. 2019) (Le et al. 2018).

This paper presents a framework for solving long-horizon temporally extended tasks with a hierarchical agent frame-

Presented at the AI-HRI Symposium at AAAI Fall Symposium Series (FSS) 2022




Figure 1: Hierarchical agent: The figure shows a scenario where all the 3 blocks are placed far away from each other. The task is to bring all the blocks close to each other. This can be broken down into four subtasks as shown in the bottom. After the low-level policy of the agent executes all four subtasks, we get the desired final configuration where all three blocks are close to each other.

work using semantic goal representations. The agent has two levels of control and the ability to easily incorporate expert supervision and intervention. The high-level policy is a symbolic planner (Alkhazraji et al. 2020) which outputs a plan in terms of sub-goals or macro-actions given an initial state and final goal state. The low-level policy is a goal-conditioned multi-task policy which is able to achieve sub-goals where these goals are specified using a semantic goal representation. The semantic goal representation is constructed using several predicate functions which define the behavior space of the agent. This representation has many benefits because it is much simpler than traditional state-based goal spaces as shown in (Akakzia et al. 2020). It is also more interpretable and easier for an expert to intervene and provide high-level feedback. For instance, given a high level goal, the planner finds a plan which can be observed by the human expert. It is easy to make small changes to the plan by adding sub-goals and changing the sequence if necessary. This framework also enables possibilities for collaboration. Due to the interpretability and ease of modifying the high-level plan, sub-tasks can be divided among agents and humans. This is not possible in other hierarchical agent frameworks where

the high-level planner is also a black-box.

We evaluate the framework using a robotic block manipulation environment. Our experiments show that this approach is able to solve different tasks by combining grasping, pushing and stacking blocks. Our contributions can be summarised as follows:

- A hierarchical agent framework where the high-level policy is a symbolic planner and the low-level policy is learned using semantic goal representations.
- Evaluation on complex long horizon robotic block manipulation tasks to show feasibility and sample efficiency
- A discussion showing the benefits of this framework in terms of interpretability and ability to interact and collaborate with humans.

Methods

In this section, we present a framework for solving long horizon temporally extended tasks. We first describe the problem statement with the environment used in our experiments. Then we describe the semantic goal representation and low-level policy training. Finally, we show how the high-level policy is obtained using the STRIPS planner (Fikes and Nilsson 1971; Alkhazraji et al. 2020) to tie everything together and solve long horizon tasks.

Problem statement

We aim to learn robotic control tasks in the Fetch Manipulation environment built on top of Mujoco (Todorov, Erez, and Tassa 2012) which consists of a robotic arm with a gripper and square blocks. The observation space consists of the arm state including positions and velocities, the gripper state, and the Cartesian positions of the blocks. The robot can pick up, push and move the blocks. We built several tasks in this domain. We have the ability to initialize the scene with different configurations, like the block positions and robot arm and gripper positions.

Semantic goal representations

We represent goals using a list of semantic predicates which are determined based on domain knowledge. In our case we consider three spatial predicates - *close*, *above*, *in-bin* and one haptic predicate - *holding*. As demonstrated by (Akakzia et al. 2020), these predicates define a much simpler behavior space instead of the traditional more complicated state space. This makes it easier to represent goals and also define a curriculum as we will discuss in the next section. And more importantly, this representation eliminates the need to write reward functions for every desired behavior.

All these predicates are binary functions applied to pairs of objects. The *close* predicate is order-invariant. $close(o_1, o_2)$ denotes whether objects (in our case blocks) o_1 and o_2 are close to each other or not. The *above* predicate is applied to all permutations of objects. $above(o_1, o_2)$ is used to denote if o_1 is above o_2 . The *in-bin* predicate is used to denote whether the block is inside the bin. Finally, *holding* is used to denote if the robot arm is holding an object using $holding(o)$. With these predicates we can form a semantic representation of the state by simply concatenating all the predicate outputs as shown in Fig 3.

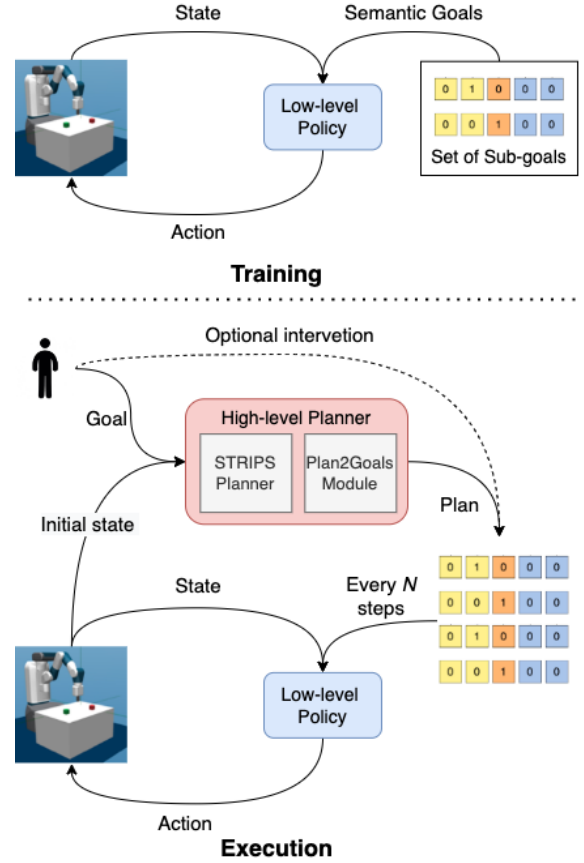


Figure 2: The low-level policy receives semantic goals and is trained to execute primitive actions in the environment to achieve the goal. The set of sub-goals is curated by a human and can also involve a curriculum. To perform high-level tasks during execution, a human can specify a goal and the high-level planner outputs a plan, a sequence a sub-goals. A human can observe this plan and intervene/modify if required. The low-level policy achieves these sub-goals sequentially to perform the task.

Training the low-level policy

The low-level policy is trained to perform several individual sub-tasks, which can eventually be used to solve longer high-level tasks. We use Hindsight experience replay (HER) (Andrychowicz et al. 2017) along with Soft-Actor critic (SAC) (Haarnoja et al. 2018) to train the goal conditioned policy. Goals are sampled from a set of configurations based on the environment where an expert can be used to optionally create a curriculum. The semantic goal representation makes it easier to do both of these things. The agent explores the environment to collect experience and updates its policy using SAC. As stated earlier, there is no need to write reward functions for each desired behavior. A reward can be generated by checking whether the current semantic configuration matches the goal configuration. The sub-goals for the two environments we use are listed in Table 1.

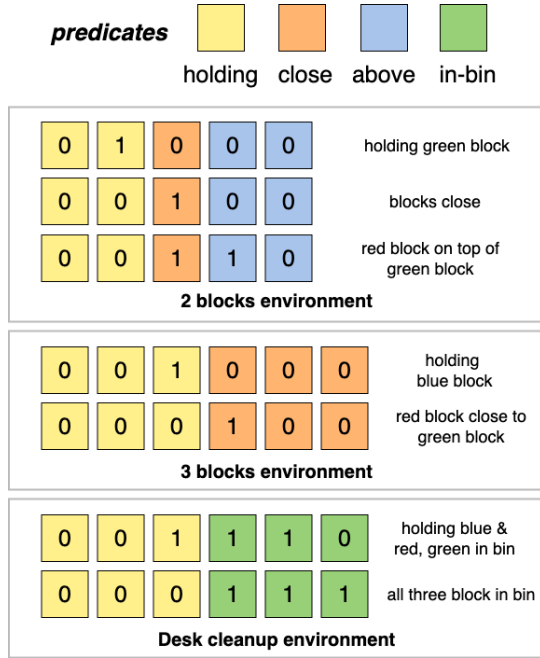


Figure 3: Semantic Predicates: This figure shows how the semantic goal representations are created using the predicates. We consider three predicates *holding*, *close*, *above* as shown. All three predicates are used in the 2 blocks environment. In the 3 blocks environment, we only use the *holding* and *close* predicates to simplify the problem. In the desk cleanup environment, we use the *holding* and *in-bin* predicates. The above figure shows examples of a few configurations.

High-level planner

We use a STRIPS planner (Fikes and Nilsson 1971; Alkharaji et al. 2020) as a high-level policy which provides sub-goals to solve the task. The STRIPS planner uses an encoding called Planning Domain Definition Language (PDDL) (Aeronautiques et al. 1998) to represent the planning task. This planner can be defined using a 5 tuple (O, P, S_i , S_g , A). Here O represents the objects of interest, P is the predicates, S_i and S_g represent the initial and goal states respectively, and A represents the action space, in this case the high level macro-actions.

The planner accepts an initial state, S_i , and goal state, S_g , represented using objects and predicates. It then uses a search algorithm like BFS, A* etc. to find a plan to reach S_g using actions A. The plan which is in PDDL is then translated to semantic goals which the low-level policy understands by *Plan2Goals* module, a function which converts the PDDL output to predicates and semantic goals. We execute the low-level policy for a fixed number of steps N , before switching the control back to the high-level sketch and execute the next sub-policy.

2 Blocks	3 Blocks	Desk Cleanup
Pick X	Pick X	Pick X
Put X near Y	Put X near Y	Put X on Table
Put X away from Y	Put X away from Y	Put X in Bin
Put X on top of Y		

Table 1: Low-level policy is trained on the above set of semantic goals. The semantic goal representation is built using the predicates as described in the previous section. For 2 blocks version, X/Y can be block of color [red, green]. For the 3 blocks version it can be of color [red, green, blue]

Experiments

Environment setup and tasks

We design two versions of the Fetch manipulation environment with 2 and 3 blocks.

2 blocks environment Here we have the robotic arm as mentioned earlier and two blocks: red and green. We consider all three predicates for this version, *close*, *above* and *holding*. We design 3 high-level tasks in this environment (1) Move blocks close: Here the task is initialized with the 2 blocks far away from each other. The goal is to bring them close to each other. (2) Move blocks apart: Here the task is initialized with blocks close to each other. The goal is to move them apart. (3) Swap blocks: Here the task is initialized with blocks on top of each other in random order. The goal is to swap the order.

3 blocks environment Here we have the robotic arm and 3 blocks: red and green and blue. For this version, we only consider 2 predicates, *close* and *holding*. We design 2 high-level tasks in this environment (1) Move blocks close: Here the task is initialized with all the 3 blocks far away from each other. The goal is to bring them close to each other. (2) Move block apart: Here the task is initialized with blocks close to each other. The goal is to move them apart.

Desk cleanup environment Here we have a robotic arm and several blocks on the desk. The desk also a bin and the blocks are places randomly on the desk. The task is to clean up the desk and place all the blocks inside the bin. We use 2 predicates here, *holding* and *in-bin*. We have 3 versions with 2, 3 and 4 blocks.

Baselines

1. *Flat semantic*: Here the agent has a single level policy but the goals are still represented using the semantic goal representations. There is no need to write reward functions for this version. 2. *Flat Continuous*: Here the goals are represented using the actual block positions of the desired configuration. The dense reward function is based on the distance between current and desired block locations and hence it is a dense reward function.

Results

We calculate task completion % for all the tasks using the fully trained agent. We train each agent for 2M steps and roll out 50 episodes using the trained policy. The values are an average of runs from three different seeds.

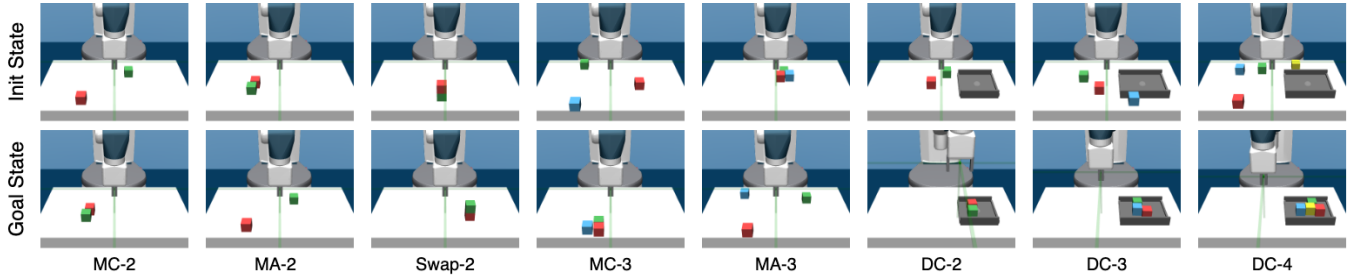


Figure 4: This figure shows all the tasks we used in our experiments. The top row shows examples of random initial states and bottom row shows the goal states. MC-2: Move 2 blocks closer, MA-2: Move 2 blocks away, Swap-2: Swap 2 stacked blocks, MC-3: Move 3 blocks closer, MA-3: Move 3 blocks away, DC-2 to DC-4: Desk clean up with 2, 3 and 4 blocks

Method	Tasks							
	MC-2	MA-2	Swap-2	MC-3	MA-3	DC-2	DC-3	DC-4
Flat Semantic	10%	80%	0%	5%	10%	30%	0%	0%
Flat Continuous	5%	10%	0%	0%	0%	0%	0%	0%
Hierarchical Semantic (Ours)	95%	100%	92%	95%	96%	94%	91%	90%

Table 2: **Task completion %** This table shows the task completion % for our experiments. The tasks names are explained in Figure 4. As seen, our method consistently outperforms all the other baselines. Our method is the only one which can consistently solve all the three tasks. We train each agent for 2M steps and roll out 50 episodes using the trained policy. The values are an average of runs from three different seeds.

Table 2 show the results for the 2 blocks, the 3 blocks and the desk cleanup environment. All the models are trained for 2M steps. As shown in the table, our method is able to solve all the eight high-level tasks. All the other baselines struggle to solve tasks. This is consistent across all the tasks. In the *move blocks apart* task represented as MA-2, shown in table 2, the flat semantic is able to learn the task but we noticed that it learns an aggressive policy where it knocks one of the blocks away from the table which not a desirable behavior. Whereas our model gently picks a block and moves it away from the other block.

To summarize, all the other baselines with and without dense reward signals fail to learn a good policy. We also performed experiments where we let the policy run for 5M steps and the baselines were unable to solve the task. This shows that using semantic predicates in the low-level policy and symbolic planner as the high-level policy truly helps to solve complex long-horizon tasks.

Discussion and Conclusion

In this paper we show that combining a symbolic planner and a low-level goal conditioned reinforcement learning policy is indeed a promising approach to build hierarchical agents. As the high-level planner and low-level policy communicate using semantic predicates, the framework is very interpretable. This also makes it easier for a human to intervene at the high-level to provide appropriate sub-goals in case the plan needs to be modified. If the environment allows multiple actors, this framework can also enable collaboration by dividing sub-tasks among them. For instance, consider the desk cleanup task with 4 blocks scattered around

the table. The planner outputs a plan with 8 sub-tasks to clean the desk. It is easy to detect independent sub-tasks - pickup red block and place in bin, pickup green block and place in bin etc. Such sequences can be assigned to a human or a second agent and the planner can re-plan which is inexpensive and fast.

There are several directions in which this framework can be extended. Currently, the human has to communicate the goals using the predicates which is already much easier than using raw states. But this could be improved by building a language interface which can translate natural language sentences to semantic goals and PDDL for the planner. With the current state space, we assumed access to predicate functions. But with more complex observation like images, one can learn these predicate functions using a small amount of labelled data. To further demonstrate the capabilities of the framework we plan to perform experiments on more complex environments, real robots and qualitative analysis using human subjects. This work is a step towards simple and interpretable hierarchical agents and we hope to build upon it.

Acknowledgments

This project was sponsored by the U.S. Army Research Laboratory under Cooperative Agreement Number W911NF2120076.

References

Aeronautiques, C.; Howe, A.; Knoblock, C.; McDermott, I. D.; Ram, A.; Veloso, M.; Weld, D.; SRI, D. W.; Barrett, A.; Christianson, D.; et al. 1998. PDDL— The Planning Domain Definition Language. *Technical Report, Tech. Rep.*

- Akakzia, A.; Colas, C.; Oudeyer, P.-Y.; Chetouani, M.; and Sigaud, O. 2020. Grounding language to autonomously-acquired skills via goal generation. *arXiv preprint arXiv:2006.07185*.
- Alkhazraji, Y.; Frorath, M.; Grützner, M.; Helmert, M.; Liebetraut, T.; Mattmüller, R.; Ortlieb, M.; Seipp, J.; Springenberg, T.; Stahl, P.; and Wülfing, J. 2020. Pyperplan. <https://doi.org/10.5281/zenodo.3700819>.
- Andrychowicz, M.; Wolski, F.; Ray, A.; Schneider, J.; Fong, R.; Welinder, P.; McGrew, B.; Tobin, J.; Pieter Abbeel, O.; and Zaremba, W. 2017. Hindsight experience replay. *Advances in neural information processing systems*, 30.
- Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, 4299–4307.
- Fikes, R. E.; and Nilsson, N. J. 1971. Strips: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3): 189–208.
- Fruit, R.; and Lazaric, A. 2017. Exploration-exploitation in mdps with options. In *Artificial Intelligence and Statistics*, 576–584. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- Jiang, Y.; Gu, S. S.; Murphy, K. P.; and Finn, C. 2019. Language as an abstraction for hierarchical deep reinforcement learning. *Advances in Neural Information Processing Systems*, 32.
- Le, H.; Jiang, N.; Agarwal, A.; Dudik, M.; Yue, Y.; and Daumé III, H. 2018. Hierarchical imitation and reinforcement learning. In *International conference on machine learning*, 2917–2926. PMLR.
- 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. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529.
- Navardi, M.; Dixit, P.; Manjunath, T.; Waytowich, N. R.; Mohsenin, T.; and Oates, T. 2022. Toward Real-World Implementation of Deep Reinforcement Learning for Vision-Based Autonomous Drone Navigation with Mission. *UMBC Student Collection*.
- Prakash, B.; Waytowich, N.; Oates, T.; and Mohsenin, T. 2021. Interactive Hierarchical Guidance using Language. *arXiv preprint arXiv:2110.04649*.
- Prakash, B.; et al. 2019. Improving Safety in Reinforcement Learning Using Model-Based Architectures and Human Intervention. In *The 32nd International Conference of the Florida Artificial Intelligence Society*. AAAI.
- Prakash, B.; et al. 2020. Guiding Safe Reinforcement Learning Policies Using Structured Language Constraints. In *SafeAI workshop Thirty-Fourth AAAI Conference on Artificial Intelligence*. AAAI.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shiri, A.; Navardi, M.; Manjunath, T.; Waytowich, N. R.; and Mohsenin, T. 2022. Efficient Language-Guided Reinforcement Learning for Resource Constrained Autonomous Systems. *IEEE Micro*.
- Shiri, A.; et al. 2021. An Energy-Efficient Hardware Accelerator for Hierarchical Deep Reinforcement Learning. In *Proceedings of IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Warnell, G.; Waytowich, N.; Lawhern, V.; and Stone, P. 2018. Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces. *AAAI Conference on Artificial Intelligence*, 1545–1553.