

Supplementary

Anonymous Author(s)

Affiliation

Address

email

1 Proof for Synthesizing Navigation Skills by Exploiting Spatial Independence

3 We now formally describe under what conditions we will resolve a manipulation option o for some
4 set of starting states Z . Consider an option o that has an associated navigation symbol σ^o to char-
5 acterize part of its initiation set I_o : $I_o = \text{Proj}(Z, S_b) \cap \sigma^o$. Then this implies that if the agent is in
6 a state that is an element of Z , and only changes the robot’s mobile base pose to be an element of
7 the navigation symbol without changing anything else, then the resulting state would be an element
8 of the initiation set of the option. We prove that if our assumptions regarding the initiation set of a
9 manipulation option are satisfied, then we can synthesize a locomotive behavior from our navigation
10 stack using our learned navigation symbol, which means we can generate the navigation stack to
11 support a specific option. Since the state is Markovian, proving for the more general case where we
12 aim to generate a navigation stack to support a manipulation plan follows from repeated applications
13 of Theorem 1, and so we omit it.

14 If a manipulation option’s o_i initiation set can be written using the definition of spatial independence
15 (Equation 1) from the current set of states Z , then sampling a location l from σ^{o_i} and synthesizing
16 and executing a path plan from the navigation stack to l from a start state in Z is sufficient for
17 enabling the robot to execute the manipulation option o_i .

18 **Theorem 1.** *If, for a starting set of states Z , the initiation set I_{o_i} for a manipulation option $o_i \in O$
19 can be characterized as in Equation 1, then a location l sampled from the associated navigation
20 symbol $l \in \sigma^{o_i}$ can be used in conjunction with a path planner to locomote the robot to a state s
21 that is within I_{o_i} the initiation set of o_i as long as there is a collision-free path.*

22 *Proof.* By our assumptions, we know that the initiation set for the manipulation option can be de-
23 composed into $I_o = \text{Proj}(Z, S_b) \cap \sigma_s^o$. We also assume that the agent starts in a state z element of Z
24 ($z \in Z$). We can then use the pose l that is sampled from the navigation symbol σ^{o_i} to synthesize
25 a navigation action n_i that starts from z and ends at location l , $n_i \in N(z, l)$ as long as there is a
26 collision free path through the environment. The effect of executing n_i from z by definition only
27 affects spatial state variables S_b , and so the resulting state is an element of $\text{Proj}(Z, S_b)$ and also an
28 element of navigation symbol σ_s^o . Therefore it the resulting state is an element of the intersection of
29 $\text{Proj}(Z, S_b)$ and σ_s^o , which is by definition the initiation set of o_i based on the Equation 1.

30

□

31 2 Simulation Experiment: Spatial Independence for Learning Symbols

32 Part of the model learning process requires identifying which factors are independence since there
33 is no a priori assumption about the structure of the initiation and effect sets of the skills. Partitioning
34 is via DBSCAN clustering [1], and the precondition classifiers are learned using a SVM [2] with
35 an RBF kernel (hyperparameters are optimized using grid search. The effect density estimation is
36 performed with a kernel density estimators [3, 4] with a Gaussian kernel, with a grid search over the
37 bandwidth.

38 To incorporate the spatial independence assumption, we use a simple augmentation to the baseline:
39 data collection occurs as before, but the spatial state variables are separated from all the other non-
40 spatial state variables before learning. We then perform model learning exactly as [5], except with
41 the separated spatial state variables already identified as an independent factor. This difference
42 represents how the spatial independence assumption can be incorporated into an existing symbol
43 learning pipeline.

44 3 Experiment: Transfer of Learned Abstractions

45 In the second set of experiments, our goal is to evaluate how AOSMs help transfer learned
46 abstractions to novel environments. For these experiments, we provided a manipulation-
47 only plan $p = \{\text{PickUp}(\text{Mug}), \text{ToggleOn}(\text{CoffeeMachine}), \text{ToggleOff}(\text{CoffeeMachine}),$
48 $\text{PutIn}(\text{Mug}, \text{CoffeeMachine}), \text{MakeCoffee}(\text{Mug}, \text{CoffeeMachine})\}$. The robot must construct the
49 navigation symbols that enable it to generate navigation behaviors that enable those actions to be
50 executed.

51 3.0.1 Approaches

52 For these set of experiments, all of the approaches perform a similar procedure. For a given scene
53 and current step of the plan o , the robot 1) uses rejection sampling to sample a pose l from the
54 associated navigation symbol σ^o 2) uses the path planner to move to location l , and 3) attempts to
55 run the manipulation option o . If the agent fails to successfully execute the manipulation option,
56 the location l is added as a negative sample to the dataset used to train σ^o ; the robot repeats these
57 steps until successful execution. When the robot is successful in executing the manipulation option,
58 location l is added as a positive sample to the dataset used to train σ^o , and the robot proceeds to the
59 next plan step. These navigation symbols are trained using Gaussian Process classifiers [6] with an
60 RBF kernel.

61 There are two important design choices when learning navigation symbols that can be chosen inde-
62 pendently of each other: 1) which spatial frame are the navigation symbols learned in, and 2) what
63 proposal distribution is used for rejection sampling. In [7], the global map frame is used as the spatial
64 frame and a random distribution for sampling, and we call this baseline **random global**. Learning
65 symbols in the map frame enables the robot to leverage a path planner to generate navigation behav-
66 iors, but it means that the robot must relearn the symbols when the scene changes. To exploit the
67 structure of object-centric skills, an object-centric spatial frame can be used to learn the symbols,
68 which the agent can transform into a map frame given a semantic map that includes object pose.
69 This enables the agent to effectively transfer learned information from one map to another. Using an
70 object-centric frame with a random sampling distribution is akin to the approach in James et al. [8],
71 which we term **random object**. However, using a uniform distribution as the proposal distribution
72 is extremely inefficient since the robot will try manipulating objects from locations extremely far
73 from the object. Kaelbling and Lozano-Pérez [9] proposed exploiting the nature of space using a
74 geometric heuristic that samples poses near the object, and so we call the baseline that uses the ge-
75 ometric heuristic for sampling poses and learning in a global map frame **heuristic global**. The final
76 approach learns in an object-centric spatial frame and uses the geometric heuristic to sample poses,
77 which to our knowledge has not been used in conjunction to learn symbols. We call this baseline
78 **heuristic object**, and it corresponds to our assumption. To give an upper-bound on performance, we
79 also evaluate an oracle, which always samples feasible manipulation locations.

80 To determine how effectively each approach enables learned abstractions to be transferred to differ-
81 ent environments, we use investigate two experimental settings: when the agent successfully finishes
82 executing the plan, 1) the scene is reset to the initial configuration and the agent retries executing the
83 plan (the single-scene setting), and 2) a new scene is chosen and the agent retries executing the plan
84 (the multi-scene setting). In the single-scene setting there is no need for transfer and the choice of
85 spatial frame does not matter. This lets us evaluate how important the chosen proposal distribution is

86 for learning navigation symbols. In the multi-scene setting, the agent must also transfer the learned
87 symbols to different scenes, which lets us evaluate how useful the choice of frame is for transfer.

88 3.0.2 Metrics

89 For each task execution in a scene, we report the cumulative total number of manipulation skills the
90 robot executed, until the plan succeeded.

91 3.0.3 Results

92 Results for the single-scene and multi-scene setting for all four approaches are in Figure ?? . The
93 approaches differ substantially between the single-scene and multi-scene setting. In the single-
94 scene setting, the heuristic sampler quickly guides the agent towards locations that afford useful
95 manipulations, when compared to sampling random locations. However, after around 15 episodes,
96 all of the approaches learn to plan in the single scene, and all approach the oracle’s performance
97 (which is just the length of the plan). However, in the multi-scene setting, although the heuristic
98 sampler with global frame starts off better than the random sampler with an object-centric frame,
99 after about 3 episodes, the object-centric frame with random sampling starts to outperform it. This
100 is because the global frame approach cannot port across different scenes, whereas object-centric
101 frames can. Therefore, we see that navigation symbols in an AOSM should be a) learned in an
102 object-centric frame to support portability to new domains, and b) learned using a sampling process
103 with geometric information included, rather than sampling at random.




Figure 1: An example demonstration of the Spot building an AOSM and using it to prepare coffee. (Left): Spot navigates around the space, identifies objects, and constructs an AOSM. (Right): With the AOSM and a manipulation-only plan, the Spot can synthesize the navigation abstractions to locomote around the environment to successfully execute the manipulation skills.

104 4 Robot Hardware Demonstration

105 Our demonstration of using an AOSM on a real robot can be seen in full detail in Figure 1. We
106 first manually drive the robot around and use an off-the-shelf SLAM implementation to generate
107 a 3D geometric map of the environment which the robot can use to navigate to 3D poses. The
108 robot then constructs a semantic map that captures the spatial pose and semantic attributes of each
109 of the relevant objects in the scene. Once the robot is equipped with a set of manipulation skills,
110 it generates an AOSM of the scene using hand-crafted navigation symbols, which enables it to
111 sample navigation poses that support successfully executing each of its manipulation skills. The
112 robot then uses a hand-specified PDDL of the coffee preparation task to generate the manipulation-
113 only plan using Fast Downward, which results in: **PickUp(WaterCup),Pour(WaterCup),**
114 **Place(WaterCup),PickUp(CoffeeGrinds),Pour(CoffeeGrinds), Place(CoffeeGrinds),** and then
115 **CloseLid(CoffeeMachine), PushButton(CoffeeMachine).** With the AOSM, the robot can syn-
116 thesize a navigation stack to support plan execution (Figure 1).

117 We time how long it takes the robot to construct an AOSM in 2 different environments. Navigating
118 the environment to observe the objects and then constructing the AOSM takes an average of 82.5
119 seconds. Executing the plan for the coffee preparation task takes on average 140 seconds. These
120 timings demonstrate the efficacy of AOSMs to enable a robot to rapidly generate the navigation
121 abstractions for supporting task execution.

122 **References**

- 123 [1] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering
124 clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.
- 125 [2] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- 126 [3] M. Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of*
127 *Mathematical Statistics*, 27(3):832, 1956.
- 128 [4] E. Parzen. On estimation of a probability density function and mode. *The Annals of Mathemat-*
129 *ical Statistics*, 33(3):1065, 1962.
- 130 [5] S. James, B. Rosman, and G. Konidaris. Autonomous learning of object-centric abstractions for
131 high-level planning. In *International Conference on Learning Representations*, 2022.
- 132 [6] H. Nickisch and C. Rasmussen. Approximations for binary gaussian process classification.
133 *Journal of Machine Learning Research*, 9(Oct):2035–2078, 2008.
- 134 [7] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez. From skills to symbols: Learning symbolic
135 representations for abstract high-level planning. *Journal of Artificial Intelligence Research*, 61:
136 215–289, 2018.
- 137 [8] S. James, B. Rosman, and G. Konidaris. Learning portable representations for high-level plan-
138 ning. In *International Conference on Machine Learning*, pages 4682–4691. PMLR, 2020.
- 139 [9] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical task and motion planning in the now. In *2011*
140 *IEEE International Conference on Robotics and Automation*, pages 1470–1477. IEEE, 2011.