Using Deep Learning to Bootstrap Abstractions for Hierarchical Robot Planning

Naman Shah
Arizona State University
Tempe, USA
namanshah@asu.edu

Siddharth Srivastava
Arizona State University
Tempe, USA
siddharths@asu.edu

ABSTRACT

This paper addresses the problem of learning abstractions that boost robot planning performance while providing strong guarantees of reliability. Although state-of-the-art hierarchical robot planning algorithms allow robots to efficiently compute long-horizon motion plans for achieving user desired tasks, these methods typically rely upon environment-dependent state and action abstractions that need to be hand-designed by experts.

We present a new approach for bootstrapping the entire hierarchical planning process. This allows us to compute abstract states and actions for new environments automatically using the critical regions predicted by a deep neural network with an auto-generated robot-specific architecture. We show that the learned abstractions can be used with a novel multi-source bi-directional hierarchical robot planning algorithm that is sound and probabilistically complete. An extensive empirical evaluation on twenty different settings using holonomic and non-holonomic robots shows that (a) our learned abstractions provide the information necessary for efficient multi-source hierarchical planning; and that (b) this approach of learning, abstractions, and planning outperforms state-of-the-art baselines by nearly a factor of ten in terms of planning time on test environments not seen during training.

KEYWORDS

Learning Abstractions for Planning; Deep Learning, Hierarchical Planning; Motion Planning; Learning for Motion Planning

ACM Reference Format:
Naman Shah and Siddharth Srivastava. 2022. Using Deep Learning to Bootstrap Abstractions for Hierarchical Robot Planning. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

Autonomous robots need to be able to efficiently compute long-horizon motion plans for achieving user desired tasks [11, 24, 27]. E.g., consider a scenario where a robot R in a household environment is tasked to reach kitchen K from its current location $B1$ (Fig. 1). State-of-the-art motion planning algorithms such as PRM [17] and RRT [22] use random sampling of low-level configurations to compute a path from the robot's current location $B1$ to its target location K. Such sampling-based methods fail to efficiently sample configurations from confined spaces such as doorways and corridors under uniform sampling [28].

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

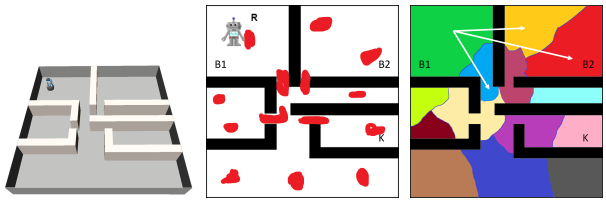

(a) (b) (c)

Figure 1: (a) An illustrative environment for a motion planning problem. The robot (R) is tasked to reach the kitchen (K). Red regions (b) show predicted critical regions in the environment. Lastly, (c) shows a projection of the computed state abstraction. Each colored cell represents an abstract state. Arrows show examples of abstract actions, defined as transitions between abstract states.

On the other hand, humans tend to reason using abstract, high-level actions. E.g., in the same scenario, we would use high-level (abstract) actions such as *"go out of the room B1"*, *"pass through the corridor"*, and *"enter the kitchen K"*. These abstract actions allow us to reason over a long-horizon easily. Currently, domain experts need to create such abstractions by hand. This limits the scope and scalability of approaches for hierarchical planning (e.g., [7, 12, 31]) to situations and domains where experts are available and able to correctly intuit the required abstractions.

This paper shows that the required abstractions can be learned by identifying regions in the environment that are important to solve motion planning problems (loosely similar to landmarks in task planning). Molina et al. [28] define such regions as *critical regions*. Critical regions are analogous to *landmarks* in automated planning but unlike landmarks, critical regions do not necessarily have to be reached to achieve the goal. Fig. 1(b) shows a few candidate critical regions for the environment in Fig. 1(a).

In this paper, we investigate two major questions: 1) Can we use deep learning to automatically generate such abstractions for new environments? And 2) how can we use them to enable safe and more efficient planning algorithms? The main contributions of this paper are: a formal foundation for hierarchical (state and action) abstractions based on the critical regions (similar to Fig. 1(c)) and a novel algorithm for on-the-fly construction of such hierarchical abstraction using predicted critical regions. Abstractions computed using learning in this manner can be difficult to use. In fact, our preliminary experiments showed that they do not help in standard single-source goal-directed search due to collision with obstacles in the environment.

An additional contribution of this work is the finding that these abstractions empower effective multi-source, multi-directional search algorithms. In general, such search algorithms can be difficult to use due to the absence of information about states that may lie close to a path to the goal. However, we found that when used with our auto-discovered abstract states, these algorithms significantly outperform existing baselines.

Our formal framework provides a way to generate sound abstractions that satisfy *downward refinement property* [1] for holonomic robots along with probabilistic completeness in the general case. Our exhaustive empirical evaluation shows that our approach outperforms state-of-the-art sampling and learning-based motion planners and requires significantly less time. This evaluation includes evaluation on a total of twenty different settings with four different robots, which include holonomic and non-holonomic robots.

The rest of the paper is organized as follows: Sec. 2 discusses some of the existing related approaches, Sec. 3 introduces some concepts required to understand our approach, Sec. 4 defines the formal framework for our algorithm, Sec. 5 presents our approach and theoretical results in detail. Finally, Sec. 6 presents our extensive empirical evaluation.

2 RELATED WORK

Much of the prior work on the topic is focused on decomposing a motion planning problem into smaller subproblems to reduce its complexity. Several approaches have been proposed that use state decomposition to reduce the complexity of a motion planning problem. *Vertical cell decomposition* [5] partitions the state space into a collection of vertical cells and computes a roadmap that passes through all of these cells. Brock and Kavraki [3] propose a hierarchical method that uses *wavefront expansion* to compute the decomposition of the state space. While these approaches establish the foundation of decomposition-based motion planning, partitions generated through such approaches are arbitrary and do not provide any guarantees of completeness.

Zhang et al. [37] use rejection sampling to reject unrelated samples to speed up SBMPs.They use reinforcement learning to learn a policy that decides to accept or reject a new sample to expand the search tree. While their approach reduces the search space to compute the path, it still needs to process samples generated from regions that are irrelevant for the current problem. On the other hand, our hierarchical approach refines abstract plans into low-level motion plans which reduces the number of unnecessary samples. TogglePRM [8] maintains roadmaps for free space and obstacle space in the configuration space to estimate the narrow passages and sample points from these narrow passages. This approach works well for environments with α-ϵ-separable passages, even though it does not compute high-level abstractions.

Multiple approaches have used statistical learning to boost motion planning. Wang et al. [35] present a comprehensive survey of methods that utilize a variety of learning methods to improve the efficiency of SBMPs. Multiple approaches discussed by Wang et al. [35] use end-to-end deep learning to learn low-level reactive policies. End-to-end approaches are attractive given if they succeed, they can compute solutions much faster than traditional approaches, but it is not exactly clear under which conditions these algorithms would succeed. Formally, these end-to-end deep learning-based approaches lack the guarantees about completeness and soundness that our approach provides. Wang et al. [35] also discuss approaches that use learning to aid sampling-based motion planning. We discuss a few of these approaches that are relevant to this work. Kurutach et al. [20] uses *InfoGAN* [6] to learn state-space partitioning for simple SE^2 robots. While their empirical evaluation shows promising results, similarly to previous decomposition-based approaches, they do not provide any proof of completeness. It is also not clear how their approach would scale to configuration spaces that had more than two dimensions. On the other hand, our approach provides formal guarantees of completeness and soundness (for holonomic robots) and scales to high-dimensional spaces.

Ichter et al. [14] and Kumar et al. [19] use a conditional variational autoencoder (CVAE) [33] to learn sampling distributions for the motion planning problems. Ichter et al. [15] use *betweenness centrality* to learn criticality score for low-level configurations. They uniformly sample a set of configurations from the environment and use configurations with higher criticality from this set to generate a roadmap. Their results show significant improvement over vanilla PRM but it is unclear how their approach would perform if the environment had regions that are important to compute motion plans yet difficult to sample under uniform sampling. On the other hand, our approach would identify such important regions to overcome these challenges. While these approaches [14, 15, 19] focus on biasing the sampling distribution towards narrow areas in the environment, our approach aims to build more general high-level abstractions for the configuration space.

Molina et al. [28] use an image-based approach to learn and infer the sampling distribution using demonstrations. They use top-view images of the environment with *critical regions* highlighted in the image to learn to identify critical regions. While they develop a method for predicting critical regions and using them with a low-level motion planner, they do not use these critical regions for learning abstractions and performing hierarchical planning. Our empirical results (Sec. 6) show that our hierarchical approach is much more effective than their non-hierarchical approach and yields significantly better performance. Additionally, their approach is also restricted to navigational problems and does not scale to configuration spaces with more than two degrees of freedom (DOFs).

Deep learning has also been used for learning heuristics for high-level symbolic planning. Shen et al. [32] use hypergraph networks for learning heuristics for symbolic planning in the form of delete-relaxation representation of the actual planning problems. Karia and Srivastava [16] learn generalizable heuristics for high-level planning without explicit action representations in symbolic logic. In contrast, we focus on learning critical regions and creating high-level abstractions along with algorithms that work with these learned high-level abstractions.

Liu et al. [23] use semantic information to bias the sampling distribution for navigational problems in partially known environments. Compared to it, our approach is not navigational problems and does not require semantic information explicitly but aims to learn such a notion in the form of critical regions. *SPARK* and *FLAME* [4] use state decomposition to store past experience and use it when queried for similar state decompositions. While their approach efficiently uses the experience from previous iterations,

it requires carefully crafted state decompositions in order to cover a large number of scenarios, whereas our approach generates state abstraction automatically using the predicted critical regions.

3 BACKGROUND

Motion Planning Problem. Let $\mathcal{X} = \mathcal{X}_{\text{free}} \cup \mathcal{X}_{\text{obs}}$ be the configuration space of a given robot [21]. Here $\mathcal{X}_{\text{free}}$ represents the set of configurations where the robot is not in collision with any obstacle and $\mathcal{X}_{\text{obs}}$ represents configurations in collision with an obstacle. Let $x_i \in \mathcal{X}_{\text{free}}$ and $x_g \in \mathcal{X}_{\text{free}}$ be the initial and goal configurations of the robot. A motion planning problem is defined as follows:

Definition 1. *A motion planning problem $\mathcal{M}$ is defined as a 4-tuple $\langle \mathcal{X}, u, x_i, x_g \rangle$. where $\mathcal{X} = \mathcal{X}_{free} \cup \mathcal{X}_{obs}$ is the configuration space and $x_i, x_g \in \mathcal{X}_{free}$ are the robot's initial and goal configurations. $u : \mathcal{X} \rightarrow \{0, 1\}$ is a collision function that determines collisions for configurations: $u(x) = 1$ iff $x \in \mathcal{X}_{obs}$.*

A solution to a motion planning problem is a collision-free trajectory $\tau : [0, 1] \rightarrow \mathcal{X}$ such that $\tau(0) = x_i$ and $\tau(1) = x_g$. We abuse the notation to define membership in a trajectory as follows: For a configuration $x \in \mathcal{X}$, $x \in \tau$ iff there a exists a $t \in [0, 1]$ such that $\tau(t) = x$. A trajectory is collision free iff $\forall x \in \tau, u(x) = 0$.

Connectivity. A pair of low-level configurations $x_i, x_j \in \mathcal{X}$ is said to be connected iff there exists a collision-free motion plan between x_i and x_j. We represent this using a connectivity function $C : \mathcal{X}_{\text{free}} \times \mathcal{X}_{\text{free}} \rightarrow \{0, 1\}$: $C(x_i, x_j) = 1$ iff x_i and x_j are connected. Intuitively, the connectivity function C represents Euclidean connectivity in $\mathcal{X}_{\text{free}}$. This is equivalent to path connectivity in configuration space for holonomic robots as each degree of freedom can be controlled independently. However, this may not be the case for non-holonomic robots as some of the motion plans may not be realizable due to their motion constraints.

We use this to define strong connectivity for the set of low-level configurations as follows:

Definition 2. *Let $\bar{\mathcal{X}} \subseteq \mathcal{X}_{free}$ be a set of configurations. $\bar{\mathcal{X}}$ is a strongly connected set iff for every pair of configurations $(x_i, x_j) \in \mathcal{X}_{free} \times \mathcal{X}_{free}$, $C(x_i, x_j) = 1$ and there exists a trajectory τ_{ij} such that $\tau_{ij}(0) = x_i$, $\tau_{ij}(1) = x_j$, and $\forall\, t \in [0, 1], \tau_{ij}(t) \in \bar{\mathcal{X}}$,*

Critical Regions. Our approach uses *critical regions (CRs)* to generate abstractions. Intuitively, critical regions are regions in the configuration space that have a high density of valid motion plans passing through them for the given class of motion planning problems. Molina et al. [28] define critical regions as follows:

Definition 3. *Given a robot R, a configuration space $\mathcal{X}$, and a class of motion planning problems M, the measure of criticality of a Lebesgue-measurable open set $r \subseteq \mathcal{X}$ is defined as $\lim_{s_n \rightarrow^+ r} \frac{f(r)}{v(s_n)}$, where $f(r)$ is the fraction of observed motion plans solving tasks from M that pass through s_n, $v(s_n)$ is the measure of s_n under a reference density (usually uniform), and $\rightarrow^+$ denotes the limit from above along any sequence $\{s_n\}$ of sets containing r ($r \subseteq s_n$, $\forall\, n$).*

Beam Search. Lowerre [25] introduced beam search as an optimization over breadth-first search (BFS) that explores the state space by expanding only a subset of nodes to compute a path from one node to another node in the graph. Beam search prunes the fringe to reduce the size of *OPEN* set. We include pseudocode for the beam search in the extended version of the paper [30].

4 FORMAL FRAMEWORK

We begin describing our formal framework with an example. Fig. 1(b) shows a set of critical regions for a given environment. Ideally, we would like to predict these critical regions and generate a state and action abstraction similar to the one shown in Fig. 1(c). The state abstraction shown in Fig. 1(c), similar to a Voronoi diagram, generates cells around each critical region such that the distance from each point in a cell to its corresponding critical regions is less than that from every other critical region. We call this structure a *region-based Voronoi diagram (RBVD)*. Each cell in this region-based Voronoi diagram is considered an abstract state and transitions between these Voronoi cells (abstract states) define abstract actions.

Let ρ be the set of critical regions for the given configuration space $\mathcal{X}$. First we introduce distance metrics d^c and d^r. Here, d^c defines distance between a low-level configuration $x \in \mathcal{X}$ and a critical region $r \in \rho$ such that $d^c(x, r) = \min_{x_i \in r} d(x, x_i)$ and d^r defines the distance between two critical regions r_1, r_2 such that the distance $d^r = \min_{x_i \in r_1, x_j \in r_2} d(x_i, x_j)$ where d is the Euclidean distance. Now, we define region-based Voronoi diagram as follows:

Definition 4. *Let $\rho = \{r_1, ..., r_k\}$ be a set of critical regions for the configuration space $\mathcal{X}$. A region-based Voronoi diagram (RBVD) is a partition $\Psi(\rho, \mathcal{X}) = \{\psi_1, ..., \psi_m\}$ of $\mathcal{X}$ such that for every $\psi_i \in \Psi$ there exists a critical regions r such that forall $x \in \psi_i$ and forall $r_j \neq r$, $d^c(x, r) \leq d^c(x, r_j)$ and each ψ_i is strongly connected.*

State Abstraction. We define abstract states as the Voronoi cells of an RBVD. Given an RBVD Ψ, labelling function $\ell : \Psi \rightarrow \mathcal{S}$ maps each cell in the RBVD to a unique abstract state $s \in \mathcal{S}$ where $|\mathcal{S}| = |\Psi|$. We use this to define the state abstraction function α as follows:

Definition 5. *Let R be the robot and $\mathcal{X} = \mathcal{X}_{free} \cup \mathcal{X}_{obs}$ be the configuration space of the robot R with set of critical regions ρ. Let $\Psi(\rho, \mathcal{X}) = \{\psi_1, ..\psi_k\}$ be an RBVD for the robot R, configuration space $\mathcal{X}$, and the set of critical regions ρ and let $\mathcal{S} = \{s_1, .., s_k\}$ be a set of high-level, abstract states. We define abstraction function $\alpha : \mathcal{X}_{free} \rightarrow \mathcal{S}$ such that $\alpha(x) = s$ where $x \in \psi$ and $\ell(\psi) = s$.*

We extend this notation to define membership in abstract states as follows: Given a configuration space $\mathcal{X} = \mathcal{X}_{\text{free}} \cup \mathcal{X}_{\text{obs}}$ and its set of abstract states $\mathcal{S}$ as defined above, a configuration $x \in \mathcal{X}_{\text{free}}$ is said to be a member of an abstract state $s \in \mathcal{S}$ (denoted $x \in s$) iff $\alpha(x) = s$. We also extend the notion of strong connectivity to abstract states as follows: An abstract state $s \in \mathcal{S}$ is strongly connected iff $\ell^{-1}(s)$ is strongly connected. We now define adjacency for Voronoi cells in a region-based Voronoi diagram as follows. Recall that C denotes Euclidean connectivity for configurations.

Definition 6. *Let ψ_i, ψ_j be Voronoi cells of an RBVD Ψ. Voronoi cells ψ_i and ψ_j are adjacent iff there exist configurations x_i, x_j such that $x_i \in \psi_i$, $x_j \in \psi_j$, $C(x_i, x_j) = 1$, and there exists a trajectory τ between x_i and x_j such that $\forall t \in [0, 1], \tau(t) \in \psi_i$ or $\tau(t) \in \psi_j$.*

We extend the above definition to define the neighborhood for an abstract state. Two abstract states $s_i, s_j \in \mathcal{S}$ are neighbors iff $\ell^{-1}(s_i)$ and $\ell^{-1}(s_j)$ are adjacent.

We define abstract actions as transitions between abstract states. Let $\mathcal{S}$ be the set of abstract states. We define the set of abstract actions $\mathcal{A}$ using $\mathcal{S}$ such that $\mathcal{A} = \{a_{ij} | \forall\ (s_i, s_j) \in \mathcal{S} \times \mathcal{S}\}$.

We now use this formulation of RBVD and state abstraction to prove the soundness of the generated abstractions .

Theorem 4.1. *Let $\mathcal{X} = \mathcal{X}_{free} \cup \mathcal{X}_{obs}$ be a configuration space and ρ be a set of critical regions for $\mathcal{X}$. Let Ψ be an RBVD for the critical regions ρ and the configuration space $\mathcal{X}$ and let $\mathcal{S}$ be the set of abstract states corresponding to Ψ with a mapping function ℓ. Let x_0 and x_g be the initial and goal configurations of a holonomic robot R. If every state $s \in \mathcal{S}$ is strongly connected and there exists a sequence of abstract states $P = \langle s_{\psi_0}, ..., s_{\psi_g} \rangle$ such that $x_0 \in s_{\psi_0}, x_g \in s_{\psi_g}$, and all consecutive states $s_{\psi_i}, s_{\psi_{i+1}} \in P$ are neighbors, then there exists a motion plan for R that reaches x_g from x_0 with a trajectory τ such that $\tau(0) = x_0$, $\tau(1) = x_g$, and $\forall x_i \in \tau, x_i \in s_{\psi_k}$ such that $s_{\psi_k} \in P$.*

Proof. For two consecutive abstract states $s_i, s_{i+1} \in P$, let $\psi_i, \psi_{i+1} \in \Psi$ be Voronoi cells such that $\ell^{-1}(s_i) = \psi_i$ and $\ell^{-1}(s_{i+1}) = \psi_{i+1}$. If s_i and s_{i+1} are neighbors, then according to Def. 6 there exists a pair of low-level configurations $x_i, x_{i+1} \in \mathcal{X}_{\text{free}}$ such that there exists a collision free trajectory between x_i to x_{i+1}. Def. 4 defines every Voronoi cell as a strongly connected set. Thus, for every low-level configuration $x_j \in s_i$, there exists a collision-free trajectory between x_j and x_i and for every low-level configuration $x_k \in s_{i+1}$, there exists a collision-free trajectory between x_k and x_{i+1}. For a holonomic robot R these trajectories should be realizable as all degrees of freedom of R can be controlled independently. This implies that there exists a motion plan for R between each pair of configurations in s_{ψ_i} and $s_{\psi_{i+1}}$.

□

Theorem 4.1 proves that the computed abstractions would be sound as well as satisfy downward refinement property for holonomic robots. The proof does not hold for non-holonomic robots as the low-level trajectories may not be realizable given their motion constraints. However, the algorithm developed below is probabilistically complete for all robots and performed well for non-holonomic robots in our empirical evaluation.

5 LEARNING ABSTRACTIONS AND PLANNING

Our approach computes hierarchical state and action abstractions using predicted critical regions and uses them efficiently for hierarchical planning. Now, we first discuss our approach for generating and using hierarchical abstractions using critical regions (Sec. 5.1) and then we discuss how we learn these critical regions (Sec. 5.2).

5.1 Generating and Using Abstractions

In this section, we describe our approach -- ***H**ierarchical **A**bstraction-guided **R**obot **P**lanner (HARP)* -- for generating abstract states and actions and using them to efficiently perform hierarchical planning. A naíve approach would be to generate a complete RBVD and then extract abstract states and actions from it. This would require iterating over all configurations in the configuration space and computing a large number of motion plans to identify executable abstract actions. This is expensive (and practically infeasible) for continuous low-level configuration spaces. Instead, we use the RBVD as an implicit concept. We generate abstractions on-the-fly by computing membership of low-level configurations in abstract states only when needed.

Algorithm 1: Hierarchical Abstraction-guided Robot Planner (HARP)

Input: Configuration space $\mathcal{X}$, a region predictor Φ, an initial configuration $x_0 \in \mathcal{X}$, goal configuration $x_g \in \mathcal{X}$, a custom heuristic h, low-level sampling-based motion planner *MP*

Output: A motion plan τ

1 $\rho \leftarrow$ predict_critical_regions(Φ, $\mathcal{X}$, x_0, x_g)
2 $\mathcal{S}, \mathcal{A} \leftarrow$ generate_state_action_abstractions(ρ, $\mathcal{X}$)
3 $s_0, s_g \leftarrow$ get_HL_state($\mathcal{S}$, ρ, x_0), get_HL_state($\mathcal{S}$, ρ, x_g)
4 $\mathcal{P} \leftarrow$ multi-source_bi-directional_beam_search($\mathcal{S}$, $\mathcal{A}$, s_0, s_g)
5 $\tau \leftarrow$ refine_path($\mathcal{P}$, *MP*)
6 $h \leftarrow$ update_heuristic(τ, $\mathcal{S}$)
7 return τ

Vanilla high-level planning using the set of all abstract actions $\mathcal{A}$ would be inefficient as it may yield plans for which low-level refinement may not exist as we do not know the applicability of these abstract actions at low-level. To overcome this challenge, we develop a hierarchical multi-source bi-directional planning algorithm that performs high-level planning from multiple abstract states. Generally, a multi-source approach would not work for robot planning because it is not clear what the intermediate states are. In this paper, we use critical regions as abstract intermediate states and utilize a multi-source search. This utilizes learned information better than single source and single direction beam search. Our high-level planner generates a set of candidate high-level plans from the abstract initial state to the abstract goal state using a custom heuristic (which is continually updated). These paths are then simultaneously refined by a low-level planner to compute a trajectory from the initial low-level configuration to the goal configuration while updating the heuristic function.

Algorithm 1 describes our approach for generating and using hierarchical abstractions. Given the configuration space $\mathcal{X}$ and initial and goal configurations (x_0 and x_g) of the robot R, HARP uses a learned DNN Φ to generate a set of critical regions ρ (Sec. 5.2 discusses how we learn this model Φ) (line 1). The remainder of Alg. 1 can be broken down into three important steps: 1) computing a set of candidate high-level plans, 2) refining candidate high-level plans into a low-level trajectory, and 3) updating the heuristic for abstract states. We now explain each of these steps in detail.

Computing High-Level Plans. To compute high-level plans that reach the goal configuration x_g from the initial configuration x_0, first we determine abstract initial and goal states s_0 and s_g corresponding to the initial and goal configurations x_0 and x_g (line 3). To do this efficiently, we store sampled points from each critical region in a k-d-tree and query it to determine the abstract state for the given low-level configurations without explicitly constructing the complete RBVD. This allows us to dynamically and efficiently determine high-level states for low-level configurations.

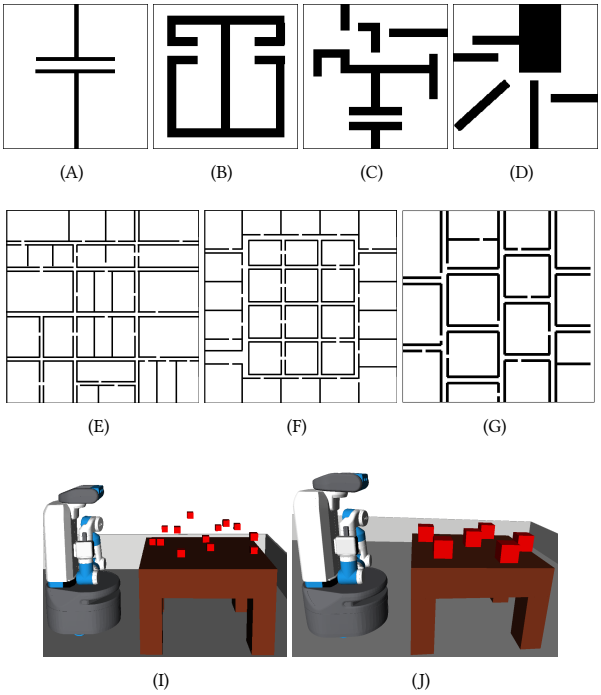

Figure 2: Test environments for our approach. Dimensions of environments (A)-(D) are 5m × 5m and dimensions of environments (E)-(G) are 25m × 25m. (A)-(G) are used for navigational problems while (H) and (I) are used for manipulation problems.

Once we determine initial and goal states s_0 and s_g, we use our high-level planner to compute a set of candidate high-level plans going from s_0 to s_g (line 4). To compute these candidate high-level plans, we develop a multi-source bi-directional variant of beam search [25] that yields multiple high-level candidate plans.

We call this *multi-source bi-directional beam search*. We include the pseudocode for vanilla and multi-source bi-directional beam search in the extended version of the paper [30].

Intuitively, this algorithm works as follows: we use a priority queue to maintain a fringe to keep track of the current state of the search. We initialize this fringe with multiple randomly sampled abstract states to allow the beam search to start from multiple sources. We select and expand these nodes in a specific order to compute paths that reach from the initial state to the goal state.

Formally, the nodes in the fringe are expanded as follows: Let n be a node in the fringe. We select the node to expand using $f(n) = g(n) + h(n)$, where $g(n)$ is the cost of the path heading to n (unit cost for each action) and $h(n)$ is a custom heuristic that is defined as follows. Let m be the parent of n and let s_n and s_m be the abstract states corresponding to the nodes n and m. The heuristic $h(n)$ is computed as:

$$h(n) = h'(s_m, s_n) + min\{h'(s_n, s_i), h'(s_n, s_g)\}$$

Here, $h'(s_1, s_2) = \epsilon_{12} d^r(r_1, r_2)$ defines the estimated distance between abstract high-level states s_1 and s_2 with corresponding critical regions r_1 and r_2 respectively. s_i is the initial abstract state and s_g is the goal abstract state.

$\epsilon_{ij} \in (0, 1]$ is a constant that accounts for imprecise abstract actions. Alg. 1 dynamically changes it to update the heuristic function (line 6) and making it more accurate. Initially, ϵ_{ij} is set to 1 for all i and j. Once a low-level trajectory τ is computed (explained later), we compute the abstraction $\bar{\tau}$ of this trajectory. For each consecutive pair of abstract states $\langle s_i, s_j \rangle$ in $\bar{\tau}$, we decrease the value of ϵ_{ij} by $\epsilon_{ij}/2$. This allows HARP to use experience from previously computed trajectories to prioritize abstract actions that have low-level refinements to compute accurate high-level plans.

After an abstract state s_j is selected from the fringe, its successors are created by applying abstract actions on it and adding these successors to the fringe. This process continues until k high-level plans are found or the fringe is empty.

Once we generate a set of candidate high-level plans from multi-source bi-directional beam search, we use a low-level planner to refine these plans into a low-level collision-free trajectory from initial low-level configuration x_0 to goal configuration x_g (line 5).

Refining High-level Plans. While any *probabilistically complete* motion planner can be used to refine the computed high-level plans into a low-level trajectory between given two configurations, we use *Learn and Link Planner (LLP)* [28] as a low-level planner in HARP (*MP* in Alg. 1) as it allows us to easily use the abstract states and their critical regions to efficiently compute motion plans. LLP is a sampling-based motion planner that initializes exploration trees rooted at N samples from the configuration space and extends these exploration trees until they connect and form a single tree. Once a single tree is formed, the planner uses Dijkstra's Algorithm [10] to compute a path from the initial state to the goal state.

In order to use LLP to refine a set of candidate high-level plans simultaneously, we first select a subset of critical regions $\bar{\rho} \subseteq \rho$ that includes all critical regions corresponding to all high-level states in candidate plans. We use this subset of critical regions $\bar{\rho}$ to provide initial samples to initialize exploration trees of LLP. In this work, we generate M samples from the set of critical regions $\bar{\rho}$ and generate the rest of the $N - M$ samples using uniform random sampling. Similarly, to expand these exploration trees, we generate a fixed number of samples from the set critical regions ρ and then continue with uniform sampling.

We use these characteristics of our algorithm to show that our approach is *probabilistic complete*.

Theorem 5.1. *If the low-level motion planner (*MP *in Alg. 1) is probabilistically complete, then HARP is probabilistically complete.*

Proof. (Sketch) While refining high-level plans to a low-level motion plan (line 5 in Alg. 1), HARP uses a fixed number of samples from the critical regions along the high-level plans to initialize a low-level motion planner. This does not reduce the set of support (regions with a non-zero probability of being sampled) of the sampling distribution being used by the motion planner. □

5.2 Learning to Predict Critical Regions

We now present our approach for learning a model Φ that predicts critical regions for a given environment. We explain how our approach is able to generate a robot-specific architecture of the

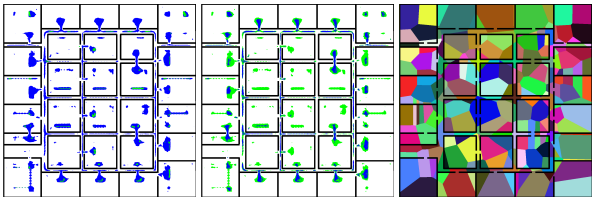

(a) (b) (c)

Figure 3: CRs and generated abstraction for a 4-DOF hinged robot. (a) and (b) Critical regions predicted by the model. Blue regions in (a) show that model predicted the robot's base link to be horizontal while green regions show that the model predicted the robot's base link to be vertical. Blue regions in (b) show that the network predicted the hinge to be closer to 180° and green regions show that the network predicted it to be closer to 90° or 270°. (c) 2D projections of state abstraction generated by our approach though our approach does not need to explicitly generate these abstractions.

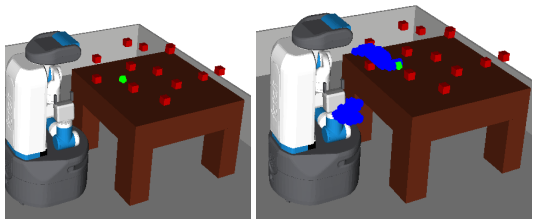

(a) (b)

Figure 4: Critical regions for arm manipulation task using an 8-DOF Fetch. The green region in (a) shows the goal location for the end effector. (b) shows the critical regions generated by the learned model. Although the network predicts CRs for all the joints, only CRs for end-effector's location are shown.

network. We include specifics about data generation and network training in the extended version of the paper [30].

Deriving Robot Specific Network Architectures. We use the standard fully convolutional UNet architecture [29] as our base network. The extended paper [30] includes network architecture for this network. We use the robot's geometry and its number of DOFs to derive a robot specific architecture as follows:

Let n be the number of DOFs of the robot and let k be the number of DOFs that are not determined by the location of the robot's end-effector in the workspace. For manipulation problems, we consider gripper of the robot as the its end-effector and for navigational problems, we consider the robot's base link as its end effector. First, we use these parameters to update the base UNet architecture. We use the last layer of the base architecture to predict critical regions for the end-effector's location and include k additional convolutional layers to predict critical regions for each k DOFs that are not determined by the location of the robot's end effector.

The input to the network is a tensor of dimension four. The size of the first three dimensions of the input tensor depends on the number of bins used to discretize the environment (which can be arbitrary). The number of channels in the input tensor is determined using the parameter n. If the robot has n DOFs, then the input tensor would have a total of $n + 1$ channels. The first channel in the input represents the occupancy matrix of the environment. It is generated by performing a raster scan of the environment. The rest of the n channels represent goal values for each DOF of the robot -- one for each DOF of the robot.

Similarly, each label is a tensor of dimension four. The size of the first three dimensions is similar to the input tensor. The number of channels in the label tensor is also computed using the robot's geometry. For a robot with k DOFs that are not determined by the location of robot's end effector in the workspace, the label tensor would have a total of $k + 1$ channels. The first channel represents critical regions for the end-effector's location in the workspace and the rest of the k channels represent critical regions for the k DOFs that are not determined by this location -- one channel for each of the k DOFs of the robot.

E.g, consider a 5-DOF hinged robot. The robot's 5 DOFs are $(x, y, z, \theta, \omega)$ where x, y, and z represent the location of the robot's base link in the workspace, θ represents the rotation of the base link, and ω represents the hinged angle. So for this robot, n would equal to 5 and k would equal to 2 as only the base rotation θ and the hinged angle ω are not determined by the location of the robot's end-effector (base link in this case) in the workspace. So according to the previous discussion, the network would contain $k = 2$ additional layers to predict critical regions for θ and ω. The input tensor would have a total of $n + 1 = 6$ channels and the label tensor would have a total of $k + 1 = 3$ channels.

6 EMPIRICAL EVALUATION

We extensively evaluate our approach in twenty different scenarios with four different robots. All experiments were conducted on a system running *Ubuntu 18.04* with *8-core i9* processor, *32 GB RAM*, and an *Nvidia 2060* GPU (our approach uses only a single core) OpenRAVE robot simulator [9]. We compare our approach with state-of-the-art motion planners such as *RRT* [22], *PRM* [17], and *BiRRT* [18]. As LLP is implemented using *Python*, we use the *Python* implementation of the baseline algorithms available at https://ompl.kavrakilab.org/ for comparison. Our training data, code, trained model, and results are available with the extended version of the paper [30] at https://aair-lab.github.io/harp.html.

3-***DOF Rectangular Robot (R).*** For the first set of experiments, the objective is to solve motion planning problems for a 3-DOF rectangular robot. The robot can move along the x and y axes and it can rotate around the z axis.

3-***DOF Non-Holonomic Rectangular Car Robot (Car).*** For the second set of experiments, we evaluated our approach with a rectangular non-holonomic robot similar to a simple car. Controls available to operate the robot were linear velocity $v \in [-0.2, 0.2]$ and the steering angle $\theta \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ while three degrees of freedom (location along x-axis, location along y-xis, and rotation around z-axis) were required to represent the robot's transformation.

4-***DOF Hinged Robot (H).*** For the third set of experiments, we used a robot with a hinge joint to evaluate our approach. The robot's

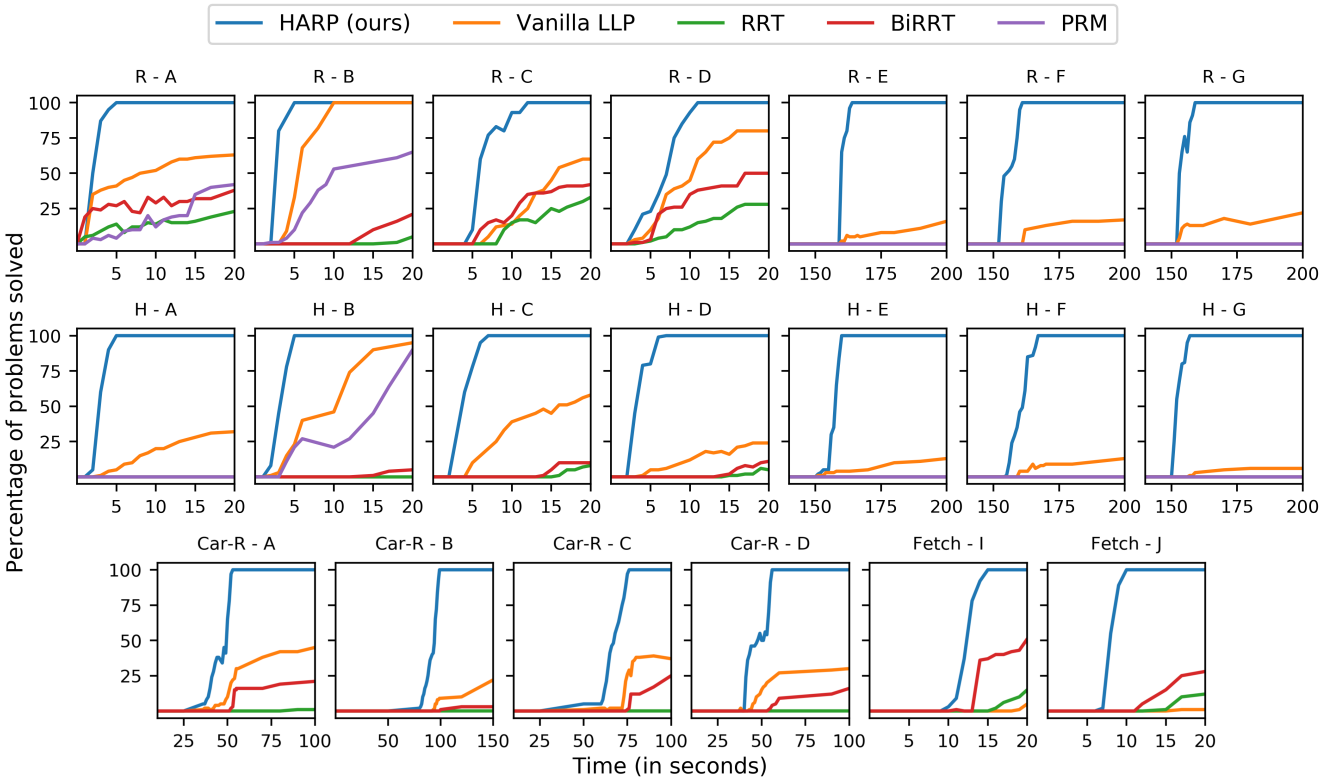

Figure 5: Each plot shows the fraction of 100 independently generated motion planning tasks solved (y-axis) in the given time (x-axis) for all the test environments and robots. The title of each subplot represents the robot and the environment. E.g., "R - A" stands for rectangular robot in environment A (Fig. 2(A)) and "H - A" stands for hinged robot in environment A.

4 DOFs are its location along x and y axes, rotation along z-axis (θ), and the hinge joint (ω) with the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$.

8-DOF Fetch Robot. For the last set of experiments, we used our approach with a mobile manipulator named Fetch [36] to perform arm manipulation. The goal of this experiment is to evaluate the scalability of our approach to robots with high degrees of freedom.

Generating Training Data. Figures 2 and 4 show the test environments (unseen by the model while training) for our system. Environments shown in Fig. 2 are inspired by the indoor office and household environments. Our training data consisted of 20 environments similar to the ones shown in Fig. 2(A)-(D) with dimensions $5m \times 5m$. We investigate the scalability of our approach by conducting experiments in environments shown in Fig. 2(E)-(G) with dimensions $25m \times 25m$ (much larger than training environments). To handle such large environments without making any changes to the DNN, we use the standard approach of sliding windows with stride equal to window-width [2, 13, 26, 34]. This crops the larger environment into pieces of the size of the training environments. Individual predictions are then combined to generate a set of critical regions for arbitrary large environments. We also evaluate the applicability of our approach to non-holonomic robots in environments shown in Fig. 2(A)-(D).

6.1 Analysis of the Results

As discussed in the introduction (Sec. 1), our objective is to show whether 1) state and action abstractions can be derived automatically and 2) whether auto-generated state and action abstractions can be efficiently used in a hierarchical planning algorithm. Additionally, we also investigate 3) does dynamically updating the heuristic function (line 6 in Alg. 1) improve Alg. 1's efficiency?

1) ***Can We Learn State and Action Abstractions?*** Our approach learns critical regions for each DOF of the robot. Fig. 3 shows critical regions predicted by our learned model for the hinged robot H. We can see that our model was able to identify critical regions in the environment such as doorways and narrow hallways. Fig. 3(a) shows critical regions for orientation of the robot's base link (captured by DOF θ). The blue regions in the figure represent the horizontal orientation of the robot and the green regions represent the vertical orientation of the robot. Fig. 3(b) shows critical regions for the hinge joint ω for the robot H. Here, blue regions show that the network predicted the hinge joint to be flat (close to 0°) and green regions represent configurations where the model predicted "L" configurations of the robot (ω close to 90° or 270°). Fig. 3 shows that our approach was able to predict the correct orientation of the robot accurately most of the time. Our approach was able to scale to robots with a high number of degrees of freedom. Fig. 4(a)

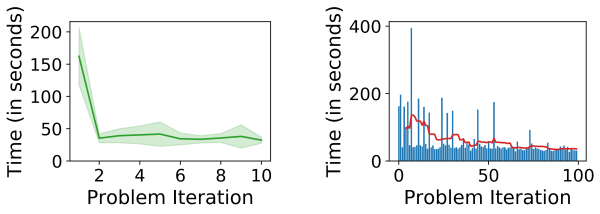

(a) (b)

Figure 6: (a) Solving 20 randomly generated problems repeatedly 110 times. x-axis show the problem iteration and y-axis shows the average time over 20 problem instances. (b) Time taken to solve 100 randomly generated problem instances. X-axis shows the problem number and y-axis shows the taken to solve each problem.

shows one of the test environments used for these experiments and Fig. 4(b) shows the predicted critical regions. This shows that our model was able to learn critical regions in the environment that can be used to generate efficient abstractions. We include similar results for other environments in the extended version [30].

Now we answer the second question on whether using abstractions to compute motion plans helps improve the planner's efficiency by qualitatively comparing our approach with a few existing sampling-based motion planners.

b) ***Can Learned Abstractions be Used Efficiently for Hierarchical Planning?*** We compare our approach against widely used SBMPs such as RRT [22], PRM [17], and BiRRT [18].

Fig. 5 shows the comparison of our approach with other sampling-based motion planners. The x-axis shows the time limit in seconds and the y-axis shows the percentage of motion planning problems solved in that time limit. For each time limit on x-axis, we randomly generate 100 new motion planning problems to thoroughly test our approach and reduce statistical inconsistencies. Fig. 5 shows that our approach significantly outperforms all of the existing sampling-based motion planners. Specifically for environments C, D, and E, uniform sampling-based approaches were not able to solve a single problem in a time threshold of 600s.

Our approach also outperforms the learning-based planner LLP [28] (Fig. 5), which uses learned critical regions but does not use state and action abstractions and does not perform hierarchical planning. This illustrates the value of learning abstractions and using them efficiently for hierarchical planning.

Similarly, we also evaluate our approach against TogglePRM [8]. TogglePRM is written in C++ and accepts only discrete SE^2 configuration space for a simple dot robot. We created discrete variants of environments shown in Fig. 2(A) and 2(B) with a total of 50176 states and compare the total number of nodes sampled. For 100 random trials, on an average our approach generated 631 ± 278 and 496 ± 175 states compared to TogglePRM which generated 4234 ± 532 and 19234 ± 4345 states for discrete variants of the environments shown in Fig. 2(A) and 2(B) respectively. Our approach was able to outperform TogglePRM since these environments do not have α-ϵ-separable passages [8].

c) ***Does Dynamically Updating Heurstic Function Improve Efficiency?*** We carried out two sets of experiments. In the first set of experiments, we generated 20 random motion planning problems and solved each problem repeatedly for 10 times while updating the heuristic function. We maintained separate copies of high-level heuristic functions for each problem. Fig. 6(a) shows the results for this set of experiments in the environment E (Fig. 2(E)) with the 4-DOF hinged robot. The x-axis shows the planning iteration and the y-axis shows the average time over randomly generated 20 problem instances. We can see how planning time reduces drastically once costs for abstract actions are updated.

In the second set of experiments, we generated 100 random pairs of initial and goal states and computed motion plans for each of them. This time, we maintained a single heuristic function across all problems and updated it after each motion planning query. Fig. 6(b) shows the result of the experiment in the environment E (Fig. 2(E)) with the 4-DOF hinged robot H. The x-axis shows the problem number and the y-axis shows the time taken by our approach to compute a solution. The red line in the plot shows the moving average of planning time. Fig. 6(b) shwows that dynamically updating the heuristic function for high-level planning helps to increase the efficiency of HARP and decrease motion planning times.

Empirical evaluation using these experiments validates our hypothesis that learning abstractions and effectively using them improves motion planning efficiency.

7 CONCLUSION

In this paper, we presented a probabilistically complete approach HARP, that uses deep learning to identify abstractions for the input configuration space. It learns state and action abstractions in a bottom-up fashion and uses them to perform efficient hierarchical robot planning. We developed a new multi-source bi-directional planning algorithm that uses learned state and action abstractions along with a custom dynamically maintained cost function to generate candidate high-level plans. A low-level motion planner refines these high-level plans into a trajectory that achieves the goal configuration from the initial configuration.

Our formal framework provides a way to generate sound abstractions that satisfy the downward refinement property for holonomic robots. Our empirical evaluation on a large variety of problem settings shows that our approach is able to significantly outperform state-of-the-art sampling and learning-based motion planners. Through our empirical evaluation, we show that our approach is robust and can be scaled to large environments and to robots that have high degrees of freedom. Our work presents a foundation for learning high-level, abstractions from low-level trajectories. Currently, our approach works for deterministic robot planning problems. We aim to extend our approach to support stochastic settings and learn abstractions for task and motion planning problems.

ACKNOWLEDGMENTS

We thank Abhyudaya Srinet for his help in implementing a primitive version of the presented work. We thank Kyle Atkinson for his help with creating test environments. This work is supported in part by the NSF under grants 1909370 and 1942856.

REFERENCES

[1] Fahiem Bacchus and Qiang Yang. 1991. The Downward Refinement Property. In *Proc. IJCAI, 1991.*

[2] Taibou Birgui Sekou, Moncef Hidane, Julien Olivier, and Hubert Cardot. 2018. Retinal Blood Vessel Segmentation Using a Fully Convolutional Network – Transfer Learning from Patch- to Image-Level. In *Proc. MLMI, 2018.*

[3] Oliver Brock and Lydia E. Kavraki. 2001. Decomposition-based Motion Planning: A Framework for Real-Time Motion Planning in High-Dimensional Configuration Spaces. In *Proc. ICRA, 2001.*

[4] Constantinos Chamzas, Zachary Kingston, Carlos Quintero-Peña, Anshumali Shrivastava, and Lydia E Kavraki. 2020. Learning Sampling Distributions Using Local 3D Workspace Decompositions for Motion Planning in High Dimensions. In *Proc. ICRA, 2021.*

[5] Bernard Chazelle. 1985. Approximation and Decomposition of Shapes. *Algorithmic and Geometric Aspects of Robotics* 1 (1985), 145–185.

[6] Xi Chen, Yan Duan, Rein Houthooft, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. In *Proc. NeurIPS, 2016.*

[7] Neil T Dantam, Zachary K Kingston, Swarat Chaudhari, and Lydia E Kavraki. 2018. An incremental constraint-based framework for task and motion planning. *The International Journal of Robotics Research* 37, 10 (2018), 1134–1151.

[8] Jory Denny and Nancy M. Amato. 2013. Toggle PRM: A Coordinated Mapping of C-Free and C-Obstacle in Arbitrary Dimension. *Algorithmic Foundations of Robotics X. Springer Tracts in Advanced Robotics (STAR)* 86, 297–312.

[9] Rosen Diankov. 2010. *Automated Construction of Robotic Manipulation Programs.* Ph.D. Dissertation. Carnegie Mellon University.

[10] Edsger W Dijkstra. 1959. A Note on Two Problems in Connexion with Graphs. *Numerische mathematik* 1, 1 (1959), 269–271.

[11] Paul Duckworth, Yiannis Gatsoulis, Ferdian Jovan, Nick Hawes, David C Hogg, and Anthony G Cohn. 2016. Unsupervised Learning of Qualitative Motion Behaviours by a Mobile Robot. In *Proc. AAMAS, 2016.*

[12] Caelan Reed Garrett, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. 2020. PDDLStream: Integrating Symbolic Planners and Blackbox Samplers via Optimistic Adaptive Planning. In *Proc. ICAPS, 2020.*

[13] Le Hou, Dimitris Samaras, Tahsin M Kurc, Yi Gao, James E Davis, and Joel H Saltz. 2016. Patch-Based Convolutional Neural Network for Whole Slide Tissue Image Classification. In *Proc. CVPR, 2016.*

[14] Brian Ichter, James Harrison, and Marco Pavone. 2018. Learning Sampling Distributions for Robot Motion Planning. In *Proc. ICRA, 2018.*

[15] Brian Ichter, Edward Schmerling, Tsang-Wei Edward Lee, and Aleksandra Faust. 2020. Learned Critical Probabilistic Roadmaps for Robotic Motion Planning. In *Proc. ICRA, 2020.*

[16] Rushang Karia and Siddharth Srivastava. 2021. Learning Generalized Relational Heuristic Metworks for Model-Agnostic Planning. In *Proc. AAAI, 2021.*

[17] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. 1996. Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE transactions on Robotics and Automation* 12, 4 (1996), 566–580.

[18] James J Kuffner and Steven M LaValle. 2000. RRT-connect: An Efficient Approach to Single-Query Path Planning. In *Proc. ICRA, 2000.*

[19] Rahul Kumar, Aditya Mandalika, Sanjiban Choudhury, and Siddhartha S. Srinivasa. 2019. LEGO: Leveraging Experience in Roadmap Generation for Sampling-Based Planning. In *Proc. IROS, 2019.*

[20] Thanard Kurutach, Aviv Tamar, Ge Yang, Stuart Russell, and Pieter Abbeel. 2018. Learning Plannable Representations with Causal InfoGAN. In *Proc. NeurIPS, 2018.*

[21] Steven . M. LaValle. 2006. *Planning Algorithms.* Cambridge University Press, Cambridge, U.K. Available at http://planning.cs.uiuc.edu/.

[22] Steven M LaValle. 1998. Rapidly-Exploring Random Trees: A New Tool for Path Planning. (1998).

[23] Katherine Liu, Maritna Stadler, and Nicholas Roy. 2020. Learned Sampling Distributions for Efficient Planning in Hybrid Geometric and Object-Level Representations. In *Proc. ICRA, 2020.*

[24] Shih-Yun Lo, Shiqi Zhang, and Peter Stone. 2018. PETLON: Planning Efficiently for Task-Level-Optimal Navigation. In *Proc. AAMAS, 2018.*

[25] Bruce T Lowerre. 1976. *The Harpy Speech Recognition System.* Carnegie Mellon University.

[26] Xin Lu, Zhe Lin, Xiaohui Shen, Radomir Mech, and James Z Wang. 2015. Deep Multi-Patch Aggregation Network for Image Style, Aesthetics, and Quality Estimation. In *Proc. ICCV, 2015.*

[27] Matteo Luperto, Luca Fochetta, and Francesco Amigoni. 2021. Exploration of Indoor Environments Predicting the Layout of Partially Observed Rooms. In *Proc. AAMAS, 2021.*

[28] Daniel Molina, Kislay Kumar, and Siddharth Srivastava. 2020. Learn and Link: Learning Critical Regions for Efficient Planning. In *Proc. ICRA, 2020.*

[29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proc. MICCAI, 2015.*

[30] Naman Shah and Siddharth Srivastava. 2022. Using Deep Learning to Bootstrap Abstractions for Hierarchical Robot Planning. arXiv:2202.00907 [cs.RO] https://arxiv.org/abs/2202.00907

[31] Naman Shah, Deepak Kala Vasudevan, Kislay Kumar, Pranav Kamojjhala, and Siddharth Srivastava. 2020. Anytime Integrated Task and Motion Policies for Stochastic Environments. In *Proc. ICRA, 2020.*

[32] William Shen, Felipe Trevizan, and Sylvie Thiébaux. 2020. Learning Domain-Independent Planning Heuristics with Hypergraph Networks. In *Proc. ICAPS, 2020.*

[33] Kihyuk Sohn, H. Lee, and Xinchen Yan. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. In *Proc. NIPS, 2015.*

[34] Chaohui Tang, Qingxin Zhu, Wenjun Wu, Wenlin Huang, Chaoqun Hong, and Xinzheng Niu. 2020. PLANET: Improved Convolutional Neural Networks with Image Enhancement for Image Classification. *Mathematical Problems in Engineering* 2020 (2020).

[35] Jiankun Wang, Tianyi Zhang, Nachuan Ma, Zhaoting Li, Han Ma, Fei Meng, and Max Q-H Meng. 2021. A survey of Learning-Based Robot Motion Planning. *IET Cyber-Systems and Robotics* (2021).

[36] Melonee Wise, Michael Ferguson, Derek King, Eric Diehr, and David Dymesich. 2016. Fetch and Freight: Standard Platforms for Service Robot Applications. In *IJCAI 2016 Workshop on Autonomous Mobile Service Robots.*

[37] Clark Zhang, Jinwook Huh, and Daniel D Lee. 2018. Learning Implicit Sampling Distributions for Motion Planning. In *Proc. IROS, 2018.*