

Planning and Execution using Inaccurate Models with Provable Guarantees

Anirudh Vemula[†], Yash Oza[†], J. Andrew Bagnell[‡] and Maxim Likhachev[†]

[†] Robotics Institute, Carnegie Mellon University

[‡] Aurora Innovation

vemula@cmu.edu

Abstract—Models used in modern planning problems to simulate outcomes of real world action executions are becoming increasingly complex, ranging from simulators that do physics-based reasoning to precomputed analytical motion primitives. However, robots operating in the real world often face situations not modeled by these models before execution. This imperfect modeling can lead to highly suboptimal or even incomplete behavior during execution. In this paper, we propose CMAX an approach for interleaving planning and execution. CMAX adapts its planning strategy online during real-world execution to account for any discrepancies in dynamics during planning, without requiring updates to the dynamics of the model. This is achieved by biasing the planner away from transitions whose dynamics are discovered to be inaccurately modeled, thereby leading to robot behavior that tries to complete the task despite having an inaccurate model. We provide provable guarantees on the completeness and efficiency of the proposed planning and execution framework under specific assumptions on the model, for both small and large state spaces. Our approach CMAX is shown to be efficient empirically in simulated robotic tasks including 4D planar pushing, and in real robotic experiments using PR2 involving a 3D pick-and-place task where the mass of the object is incorrectly modeled, and a 7D arm planning task where one of the joints is not operational leading to discrepancy in dynamics. The video of our physical robot experiments can be found at <https://youtu.be/eQmAeWihjO8>.

I. INTRODUCTION

Modern robotic planning approaches involve use of models that tend to be sophisticated and complex. These models are used to simulate the dynamics of the real world and foresee the outcomes of actions executed. From using fast analytical solvers to generate motion primitives on-the-fly [9] to simulators that do reasoning based on physics, and optimization to resolve contacts [30], these models are getting better at modeling the dynamics of the real world. However, real world robotic tasks are rife with situations that cannot be predicted and therefore, modeled before execution. Thus, we need a planning approach that can use potentially inaccurate models and still complete the task.

For example, consider the task depicted in Figure 1 (left) where a robotic arm needs to pick an object and place it at a goal location. Without knowledge of the mass of the object, the model can be inaccurate in simulating the dynamics. If the

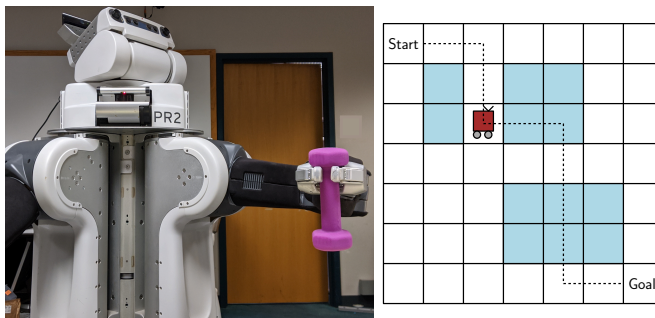


Fig. 1: (left) PR2 executing a pick-and-place task with a heavy object that is modeled as light, resulting in hitting joint torque limits during execution. (right) Mobile robot navigating a gridworld with icy states, where the robot slips, that are not modeled as icy resulting in discrepancy in dynamics.

object is modeled as light, the planned path would pick it to a certain height before placing it at the goal location. However, if the object is heavy in the real world, like in Figure 1 (left), this plan cannot be executed as the joint torque limits are reached and the arm cannot move higher. Thus, by using the inaccurate model for planning, the arm is stuck and cannot reach the goal. Figure 1 (right) presents another simple scenario where a mobile robot is navigating a gridworld containing icy states, where the robot slips, i.e. if the robot tries to go right or left in an icy state, it will move two cells rather than one cell in that direction. However, the model used for planning does not model the icy states and hence, cannot simulate the real world dynamics correctly. This can lead to highly suboptimal paths or sometimes even inability to reach the goal, when using such a model for planning.

A typical solution to this problem is to update the dynamics of the model and replan [29]. However, this is often impossible in real world planning problems where we use models that are complex and in some cases obtained from expensive computation that is done offline before execution [12]. The dynamics of these models cannot be changed online arbitrarily without deteriorating their simulation capabilities in other scenarios and sacrificing real-time execution. In addition, this solution might require us to have the knowledge of what part of the model dynamics is inaccurate and how to correct it. Going back to the pick-and-place example in Figure 1, to update the model we need to first identify that the modeled mass is incorrect and then estimate the true mass to correct the dynamics of the model. Both of these steps require specialized non-trivial implementations. Finally, in the case of models

A blog post summarizing this work can found at <https://vvanirudh.github.io/blog/cmax/>, and the full version of the paper (including appendix and acknowledgements) can be found at <https://arxiv.org/abs/2003.04394>

This work was in part supported by ONR grant N00014-18-1-2775

that *can* be updated online efficiently, it might still not be possible to model the true dynamics without an unreasonably large number of online executions because the true dynamics are often very complex, e.g. modeling cooperative navigation dynamics in human crowds [31]. The above aspects make the solution of updating model dynamics online undesirable in real world robotic tasks, where we are interested in completing the task and *not* in modeling the dynamics accurately.

In this work, we present an alternative approach CMAX for interleaving planning and execution that does not require updating the dynamics of the model. Instead during execution, whenever we discover an action where the dynamics differ between the real world and the model, we update the cost function to penalize executing such state-action pairs in the future. This biases the planner to replan paths that do not consist of such state-action pairs, and thereby avoid regions of state-action space where the dynamics are known to differ. Based on this idea, we present algorithms for both small state spaces, where we can do exact planning, and large state spaces, including continuous state spaces, where we resort to function approximation to update the cost function and to maintain cost-to-go estimates. Our framework CMAX comes with provable guarantees on reaching the goal, without any resets, under specific assumptions on the model. The proposed algorithms are tested on a range of tasks including simulated 4D planar pushing as well as physical robot 3D pick-and-place task where the mass of the object is incorrectly modeled, and 7D arm planning tasks when one of the joints is not operational, leading to discrepancy in dynamics.

II. PRELIMINARIES

We are interested in the deterministic shortest path problem represented by the tuple $M = (\mathbb{S}, \mathbb{A}, \mathbb{G}, f, c)$ where \mathbb{S} denotes the state space, \mathbb{A} denotes the action space, $\mathbb{G} \subseteq \mathbb{S}$ is the non-empty set of goal states we are interested in reaching, $f : \mathbb{S} \times \mathbb{A} \rightarrow \mathbb{S}$ denotes the deterministic dynamics governing the transition to next state given current state and action, and $c : \mathbb{S} \times \mathbb{A} \rightarrow [0, 1]$ is the cost function. For the purposes of this work, we will focus on small discrete action spaces, bounded costs lying between 0 and 1¹, and a cost-free termination goal state i.e. for all $g \in \mathbb{G}$, we have $c(g, a) = 0$ and $f(g, a) = g$ for all actions $a \in \mathbb{A}$. The objective of the shortest path problem is to find the least-cost path from any given start state $s_0 \in \mathbb{S}$ to any goal state $g \in \mathbb{G}$ in M . We assume that there exists at least one path from each state $s \in \mathbb{S}$ to one of the goal states $g \in \mathbb{G}$ in M , and that the cost of any transition starting from a non-goal state is positive i.e. $c(s, a) > 0$ for all $s \in \mathbb{S} \setminus \mathbb{G}, a \in \mathbb{A}$. These assumptions are typical for analysis in deterministic shortest path problems [6]. We use $V(s)$ to denote the cost-to-go estimate of any state $s \in \mathbb{S}$ and $V^*(s)$ to denote the optimal cost-to-go. From dynamic programming literature [6], we know that the optimal cost-to-go satisfies the Bellman optimality condition $V^*(s) = \min_{a \in \mathbb{A}} [c(s, a) + V^*(f(s, a))]$. A cost-to-go estimate V is called admissible if

it underestimates the optimal cost-to-go $V(s) \leq V^*(s)$ for all $s \in \mathbb{S}$, and is called consistent if it satisfies the condition that for any state-action pair $(s, a), s \notin \mathbb{G}, V(s) \leq c(s, a) + V(f(s, a))$, and $V(g) = 0$ for all $g \in \mathbb{G}$.

In this work, we assume that the exact dynamics are initially unknown to the robot, and can only be discovered through executions. Thus, instead of offline planning methods, we need online methods that interleave planning with action execution. Specifically, we focus on the online real-time planning setting where the robot does not have access to resets, and the robot has to interleave planning and execution to ensure real-time operation. This is similar to the classical real-time search setting considered by works like LRTA* [22], RTAA* [20], RTDP [4] and several others. An important aspect of these approaches is that the robot can only perform a fixed amount of computation for planning, independent of the size of state space, before it has to execute an action.

III. PROBLEM SETUP

Consider the problem of a robot acting to find a least-cost path to a goal in an environment represented by the tuple $M = (\mathbb{S}, \mathbb{A}, \mathbb{G}, f, c)$ with unknown deterministic dynamics f and known cost function c . The robot gathers knowledge of the dynamics over a single trajectory in the environment, and does not have access to any resets, ruling out any episodic approach. This is an extremely challenging setting as the robot has to reason about whether to exploit its current knowledge of the dynamics to act near-optimally or to explore to gain more knowledge of the dynamics, possibly at the expense of suboptimality.

We assume that the agent has access to an approximate model, $\hat{M} = (\mathbb{S}, \mathbb{A}, \mathbb{G}, \hat{f}, c)$, that it can use to simulate the outcome of its actions and use for planning. In our motivating gridworld example (Figure 1 right), this model represents a grid with no icy states, so the dynamics \hat{f} moves the robot to the next cell based on the executed action without any slip. However, the real environment contains icy states resulting in dynamics f that differ on state-action pairs where the state is icy. For the remainder of this paper, we will refer to such state-action pairs where f and \hat{f} differ as “incorrect” state-action pairs, and use the notation $\mathcal{X} \subseteq \mathbb{S} \times \mathbb{A}$ to denote the set of “incorrect” state-action pairs, i.e. $f(s, a) \neq \hat{f}(s, a)$ for all $(s, a) \in \mathcal{X}$. The objective is for the robot to reach a goal state from a given start state, despite using an inaccurate model for planning, while minimizing the cost incurred and ensuring real-time execution.

IV. APPROACH

Existing planning and learning approaches try to learn a very good approximation of M from scratch through online executions [17, 7, 15, 10], or update the dynamics of model \hat{M} so that it approximates M well [1, 14, 27]. In this work, we propose an approach CMAX that uses the inaccurate model \hat{M} online *without* updating its dynamics, and is provably guaranteed to complete the task. In a nutshell, instead of learning a new dynamics model from scratch or updating

¹Any bounded non-negative cost can be scaled to fit this assumption

the dynamics of existing model, CMAX maintains a running estimate of the set \mathcal{X}_t consisting of all state-action pairs that have been executed and have been discovered to be incorrect until timestep t . Using the set \mathcal{X}_t , we update the cost function to bias the planner to plan future paths that avoid state-action pairs that are known to be incorrect. It is important to note that the challenge of dealing with exploration-exploitation dilemma online still exists, as we do not know the set of state-action pairs \mathcal{X} where the dynamics differ ahead of online execution. A similar approach was proposed in Jiang [14] for the episodic setting where the robot had access to resets, and for small state spaces where we could perform full state space planning. CMAX extends it to the significantly more challenging online real-time setting and we present a practical algorithm for large state spaces.

A. Penalized Model

We formalize our approach as follows: Given a model \hat{M} and a set $\mathcal{X} \subseteq \mathbb{S} \times \mathbb{A}$ consisting of state-action pairs that have been discovered to be incorrect so far, define the penalized model $\tilde{M}_{\mathcal{X}}$ as:

Definition 4.1 (Penalized Model): The penalized model $\tilde{M}_{\mathcal{X}} = (\mathbb{S}, \mathbb{A}, \mathbb{G}, \hat{f}, \tilde{c}_{\mathcal{X}})$ has the same state space, action space, set of goals, and dynamics as \hat{M} . The cost function $\tilde{c}_{\mathcal{X}}$ though is defined as $\tilde{c}_{\mathcal{X}}(s, a) = |\mathbb{S}|$ if $(s, a) \in \mathcal{X}$, else $\tilde{c}_{\mathcal{X}}(s, a) = c(s, a)$.²

Intuitively, the penalized model $\tilde{M}_{\mathcal{X}}$ has a very high cost for any transition where the dynamics differ, i.e. $(s, a) \in \mathcal{X}$, and the same cost as the model \hat{M} otherwise. More specifically, the cost is inflated to the size of the statespace, which is the maximum cost of a path that visits all states³ (remember, that our cost is normalized to lie within 0 and 1.) This biases the planner to “explore” all other state-action pairs that are not yet known to be incorrect before it plans a path through an incorrect state-action pair. In the next section, we will describe how we use the penalized model $\tilde{M}_{\mathcal{X}}$ for real-time planning.

B. Limited-Expansion Search for Planning

During online execution, the robot has to constantly plan the next action to execute from its current state in real-time. This forces the robot to use a fixed amount of computation for planning before it has to execute the best action found so far. In this work, we use a real-time search method that is adapted from RTAA* proposed by Koenig and Likhachev [20].

The planner is summarized in Algorithm 1. At any timestep t , given the current penalized model $\tilde{M}_{\mathcal{X}_t}$ and the current state s_t , the planner constructs a lookahead search tree using at most K state expansions. We obtain the successors of any expanded state and the cost of any state-action pair using the penalized model $\tilde{M}_{\mathcal{X}_t}$. After expanding K states, it finds the best state s_{best} among the leaves of the search tree that has the least sum of cost-to-come from s_t and cost-to-go to a goal state (line 18 in Algorithm 1). The best action to execute in the current state s_t is chosen to be the first action on the

Algorithm 1 Limited-Expansion Search based on RTAA*[20]

```

1: function SEARCH( $s, \tilde{M}_{\mathcal{X}}, V, K$ )
2:   Initialize  $g(s) \leftarrow 0$ 
3:   Initialize min-priority open list  $O$ , and closed list  $C$ 
4:   Add  $s$  to open list  $O$  with priority  $g(s) + V(s)$ 
5:   for  $i = 1, 2, \dots, K$  do
6:     Pop  $s_i$  from open list  $O$ 
7:     If  $s_i \in \mathbb{G}$ , then  $s_{\text{best}} \leftarrow s_i$  and move to Line 19
8:     for  $a \in \mathbb{A}$  do  $\triangleright$  Expanding state  $s_i$ 
9:       Get successor  $s' = \hat{f}(s_i, a)$ 
10:      If  $s' \in C$ , continue to next action
11:      if  $s' \in O$  and  $g(s') > g(s_i) + \tilde{c}_{\mathcal{X}}(s_i, a)$  then
12:        Update  $g(s') \leftarrow g(s_i) + \tilde{c}_{\mathcal{X}}(s_i, a)$ 
13:        Reorder open list  $O$ 
14:      else if  $s' \notin O$  then
15:        Set  $g(s') \leftarrow g(s_i) + \tilde{c}_{\mathcal{X}}(s_i, a)$ 
16:        Add  $s'$  to  $O$  with priority  $g(s') + V(s')$ 
17:      Add  $s_i$  to the closed list  $C$ 
18:   Pop  $s_{\text{best}}$  from open list  $O$ 
19:   for  $s' \in C$  do
20:     Update  $V(s') \leftarrow g(s_{\text{best}}) + V(s_{\text{best}}) - g(s')$ 
21:   Backtrack from  $s_{\text{best}}$  to  $s$ , and set  $a_{\text{best}}$  as the first
   action on path from  $s$  to  $s_{\text{best}}$ 
   return  $a_{\text{best}}$ 

```

path from s_t to s_{best} in the search tree and the cost-to-go estimates of all expanded states are updated as: $V(s_{\text{expanded}}) = g(s_{\text{best}}) + V(s_{\text{best}}) - g(s_{\text{expanded}})$, where $g(s)$ is the cost-to-come from s_t for any state s in the search tree. The amount of computation used to compute the best action for the current state is bounded as a factor of the number of expansions K in the search tree. Thus, we can bound the planning time and ensure real-time operation for our robot.

C. Warm Up: Small State Spaces

In this section, we will present an algorithm that is applicable for small discrete state spaces where it is feasible to maintain cost-to-go estimates for all states $s \in \mathbb{S}$ using a tabular representation, and we can maintain a running set \mathcal{X}_t containing all the discovered incorrect state-action pairs so far, without resorting to function approximation. The algorithm⁴ is shown in Algorithm 2. Intuitively, Algorithm 2 maintains a running set of incorrect state-action pairs \mathcal{X}_t , updates the set whenever it encounters an incorrect state-action pair, and recomputes the penalized model $\tilde{M}_{\mathcal{X}_t}$. Crucially, the algorithm never updates the dynamics of the model \hat{M} , and only updates the cost function according to Definition 4.1. In order to prove completeness, we assume the following:

Assumption 4.1: Given a penalized model $\tilde{M}_{\mathcal{X}_t}$ and the current state s_t at any timestep t , there always exists at least one path from s_t to a goal state that *does not contain* any

²This is similar to the notion of penalized MDP, introduced in Jiang [14]

³Hence, the name CMAX for our approach

⁴A similar algorithm in the episodic setting with full state space planning is presented in Jiang [14]

Algorithm 2 CMAX – Small State Spaces

```
1: Initialize  $\hat{M}_1 \leftarrow \hat{M}$ ,  $\mathcal{X}_1 \leftarrow \{\}$ , start state  $s_1 \in \mathbb{S}$ , cost-to-go estimates  $V$ , number of expansions  $K$ ,  $t \leftarrow 1$ 
2: while  $s_t \notin \mathbb{G}$  do
3:   Get  $a_t = \text{SEARCH}(s_t, \hat{M}_t, V, K)$ 
4:   Execute  $a_t$  in environment  $M$  to get  $s_{t+1} = f(s_t, a_t)$ 
5:   if  $s_{t+1} \neq \hat{f}(s_t, a_t)$  then
6:     Add  $(s_t, a_t)$  to the set :  $\mathcal{X}_{t+1} \leftarrow \mathcal{X}_t \cup \{(s_t, a_t)\}$ 
7:     Update the penalized model :  $\hat{M}_{t+1} \leftarrow \hat{M}_{\mathcal{X}_{t+1}}$ 
8:   else
9:      $\mathcal{X}_{t+1} \leftarrow \mathcal{X}_t$ ,  $\hat{M}_{t+1} \leftarrow \hat{M}_t$ 
10:   $t \leftarrow t + 1$ 
```

state-action pairs (s, a) that are known to be incorrect, i.e. $(s, a) \in \mathcal{X}_t$.⁵

Under this assumption, we can show the following guarantee for Algorithm 2:

Theorem 4.1: Assume Assumption 4.1 holds then, if \mathcal{X} denotes the set consisting of all incorrect state-action pairs, and the initial cost-to-go estimates used are admissible and consistent, then using Algorithm 2 the robot is guaranteed to reach a goal state in at most $|\mathbb{S}|^2$ timesteps. Furthermore, if we allow for $K = |\mathbb{S}|$ expansions, then we can guarantee that the robot will reach a goal state in at most $|\mathbb{S}|(|\mathcal{X}| + 1)$ timesteps.

Proof of the above theorem is given in Appendix A. The above theorem establishes that using Algorithm 2, the robot is guaranteed to reach a goal state under Assumption 4.1. In practice, we observe that the number of timesteps to reach a goal has a smaller dependence on the size of state space than the worst-case bound, especially if Algorithm 2 starts with cost-to-go estimates that are reasonably accurate for the initial model \hat{M} .

D. Large State Spaces

In large state spaces, it is infeasible to maintain cost-to-go estimates for all states $s \in \mathbb{S}$ using a tabular representation and maintain a running estimate of the set \mathcal{X}_t , as both could be very large in size. Thus, we will need to resort to function approximations for both cost-to-go estimates and the set \mathcal{X}_t .

We will assume existence of a fixed distance metric $d : \mathbb{S} \times \mathbb{S} \rightarrow \mathbb{R}^+ \cup \{0\}$, and that \mathbb{S} is bounded under this metric. We relax the definition of \mathcal{X} using the distance metric d as follows: Define any state-action pair $(s, a) \in \mathcal{X}^\xi$ to be ξ -incorrect if $d(f(s, a), \hat{f}(s, a)) > \xi$ where $\xi \geq 0$. We assume that there is an underlying path following controller that is used to execute our plan and can deal with discrepancies smaller than ξ . Thus, we allow for small discrepancies in our approximate model \hat{M} that can be resolved using a low-level controller.

Our algorithm for large state spaces is presented in Algorithm 3. The main idea of the algorithm is to “cover” the set

\mathcal{X}^ξ using hyperspheres in $\mathbb{S} \times \mathbb{A}$. Since the action space \mathbb{A} is a discrete set, we maintain separate sets of hyperspheres for each action $a \in \mathbb{A}$. Whenever the agent encounters an incorrect state-action pair $(s, a) \in \mathcal{X}^\xi$, it places a hypersphere at s corresponding to action a whose radius (as measured by the metric d) is given by $\delta > 0$, a domain-dependent constant. We inflate the cost of a state-action pair (s, a) , according to Definition 4.1, if s lies inside any hypersphere corresponding to action a . In practice, this is implemented by constructing separate KD-Trees in state space \mathbb{S} for each action $a \in \mathbb{A}$ to enable efficient lookup.

After executing the action and placing a hypersphere if a discrepancy in dynamics was observed, the function approximation for cost-to-go is updated iteratively as follows (Line 15 to Line 17): Sample a batch of states from the buffer of previously visited states with replacement, construct a lookahead tree for each state in the batch (through parallel jobs) to obtain all states on the closed list and their corresponding cost-to-go updates using Algorithm 1, and finally update the parameters of the cost-to-go function approximator to minimize the mean squared loss $\mathcal{L}(V_\theta, \mathbb{X}) = \frac{1}{2|\mathbb{X}|} \sum_{(s, V(s)) \in \mathbb{X}} (V(s) - V_\theta(s))^2$ for all the expanded states through a gradient descent step (Line 17).

Observe that, similar to Algorithm 2, we do not update the dynamics \hat{f} of the model, and only update the cost function according to Definition 4.1. However, unlike Algorithm 2, we do not explicitly maintain a set of incorrect state-action pairs but maintain it implicitly through hyperspheres. By using hyperspheres, we obtain local generalization and increase the cost of all the state-action pairs inside a hypersphere. In addition, unlike Algorithm 2, we update cost-to-go estimates of not only the expanded states in the lookahead tree obtained from current state s_t , but also from previously visited states. This ensures that the function approximation used for maintaining cost-to-go estimates does not deteriorate for states that were previously visited, and potentially help in generalization.

We can provide a guarantee on the completeness of Algorithm 3 by assuming the following:

Assumption 4.2: Given a penalized model $\tilde{M}_{\mathcal{X}_t^\xi}$ and the current state s_t at any timestep t during execution, there always exists at least one path from s_t to a goal state that is at least δ distance away from any state-action pair (s, a) that is known to be ξ -incorrect, i.e. $(s, a) \in \mathcal{X}_t^\xi$.

The above assumption has two components: the first one relaxes Assumption 4.1 to accommodate the notion of ξ -incorrectness, and the second one states that, unlike Assumption 4.1, there exists a path that not only does not contain any state-action pairs that are known to be ξ -incorrect, but also that any state-action pair on the path is at least δ distance, as measured by the metric d , away from any state-action pair that is known to be ξ -incorrect. The second component makes this assumption stronger. However, it can lead to substantial speedups in the time taken to reach a goal as we can place hyperspheres of radius δ to quickly “cover” the ξ -incorrect set.

Algorithm 3 employs approximate planning by using a function approximator for cost-to-go estimates and performing

⁵This assumption is less restrictive than the assumption that there exists at least one path from the current state to a goal that does not contain any state-action pairs (s, a) that are incorrect i.e. $(s, a) \in \mathcal{X}$

Algorithm 3 CMAX – Large State Spaces

- 1: Initialize $\hat{M}_1 \leftarrow \hat{M}$, Cost-to-go function approximation V_{θ_1} , Set of hyperspheres $\mathcal{X}_1^\xi \leftarrow \{\}$, Start state s_1 , Number of planning updates N , Batch size B , Buffer \mathcal{D} , Number of expansions K , Learning rate η , $t \leftarrow 1$, Radius of hypersphere δ , Discrepancy threshold ξ
- 2: **while** $s_t \notin \mathbb{G}$ **do**
- 3: Get $a_t \leftarrow \text{SEARCH}(s_t, \hat{M}_t, V_{\theta_t}, K)$
- 4: Execute a_t in environment M to get $s_{t+1} \leftarrow f(s_t, a_t)$
- 5: **if** $d(s_{t+1}, \hat{f}(s_t, a_t)) > \xi$ **then**
- 6: Add $\mathcal{X}_{t+1}^\xi \leftarrow \mathcal{X}_t^\xi \cup \{\text{sphere}(s_t, a_t, \delta)\}$
- 7: **else**
- 8: $\mathcal{X}_{t+1}^\xi \leftarrow \mathcal{X}_t^\xi$
- 9: Update $\hat{M}_{t+1} \leftarrow \tilde{M}_{\mathcal{X}_{t+1}^\xi}$
- 10: Add s_t to buffer \mathcal{D}
- 11: Update $V_{\theta_{t+1}} \leftarrow \text{UPDATE}(s_t, \hat{M}_{t+1}, V_{\theta_t}, \mathcal{D})$
- 12: $t \leftarrow t + 1$
- 13: **function** $\text{UPDATE}(s, \hat{M}, V_\theta, \mathcal{D})$
- 14: **for** $n = 1, \dots, N$ **do**
- 15: Sample batch of B states S_n from buffer \mathcal{D} with replacement
- 16: Call $\text{SEARCH}(s_i, \hat{M}, V_\theta, K)$ for each $s_i \in S_n$ to get all states on closed list s'_i and their corresponding cost-to-go updates $V(s'_i)$ and construct the training set $\mathbb{X}_n = \{(s'_i, V(s'_i))\}$
- 17: Update: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(V_\theta, \mathbb{X}_n)$
- return** V_θ

batch updates to fit the approximator. This is necessary as the state space is large, and maintaining tabular cost-to-go estimates for each state is expensive in memory and would take a large number of timesteps to update them in practice. However, for ease of analysis, we will assume that we do exact updates and maintain tabular cost-to-go estimates like Algorithm 2. Then, we can show the following guarantee:

Theorem 4.2: Assume Assumption 4.2 holds then, if \mathcal{X}^ξ denotes the set of all ξ -incorrect state-action pairs, and the initial cost-to-go estimates are admissible and consistent, then using Algorithm 3 with exact updates and tabular representation for cost-to-go estimates, the robot is guaranteed to reach a goal state in at most $|\mathbb{S}|^2$ timesteps. Furthermore, if we allow for $K = |\mathbb{S}|$ expansions, then we can guarantee that the robot will reach a goal state in at most $|\mathbb{S}|(\mathcal{C}(\delta) + 1)$ timesteps, where $\mathcal{C}(\delta)$ is the covering number of the set \mathcal{X}^ξ .

Proof of the above theorem is given in Appendix B. The above theorem states that, using Algorithm 3, the robot is guaranteed to reach a goal state, if the initial cost-to-go estimates are admissible and consistent. The theorem also provides a stronger guarantee that the number of timesteps to the goal has a dependence on the covering number, if we do $|\mathbb{S}|$ number of expansions at each timestep. Covering number $\mathcal{C}(\delta)$ of a set A is formally defined as the size of the set B of state-action pairs (s, a) such that $A \subseteq \bigcup_{(s,a) \in B} \text{sphere}(s, a, \delta)$. Note that the covering number $\mathcal{C}(\delta)$ is typically much smaller than the

	Accurate Model		Inaccurate Model	
	Steps	% Success	Steps	% Success
CMAX	63 ± 22	90%	192 ± 40	80%
Q-Learning	34 ± 5	90%	441 ± 100	45%
Model NN	62 ± 26	90%	348 ± 82	15%
Model KNN	106 ± 34	95%	533 ± 118	50%
Plan with Acc. Model	63 ± 22	90%	364 ± 53	85%

TABLE I: Results for the simulated 4D planar pushing task. First column corresponds to the case when the environment has no obstacles, and the model is accurate. Second column corresponds to when the environment has static obstacles, and model (with no obstacles) is inaccurate. Each entry in the Steps subcolumn is obtained using 20 random start and goal locations, and we present mean and standard error of number of timesteps it takes the robot to reach the goal among successful trials. The % success subcolumn indicates percentage of successful trials where the robot reached the goal in less than 1000 timesteps. The last row corresponds to using the planner with an accurate model (the same as the environment).

size of the set \mathcal{X}^ξ . Although performing $|\mathbb{S}|$ expansions at each timestep is infeasible in large state spaces with real-time constraints, it is useful to note that we achieve speedup from adding hyperspheres of radius δ . Importantly, the efficiency of the Algorithm 3 degrades gracefully with decreasing δ and reduces to the bound presented in Theorem 4.1, if only Assumption 4.1 holds. Similar to the worst-case bounds presented in Theorem 4.1, the number of timesteps it takes for the robot to reach a goal state, in practice as shown in our experiments, has a much smaller dependence on size of state space if we start with cost-to-go estimates that are reasonably accurate for the initial model \hat{M} , and use cost-to-go function approximation as we do in Algorithm 3.

V. EXPERIMENTS

We test the applicability and efficiency of our approach CMAX on a range of robotic tasks across simulation and real-world experiments.⁶ In simulated experiments, we record the mean and standard error for the number of timesteps taken by the robot to reach the goal emphasizing the performance of CMAX. For physical robot experiments, we present real-time execution statistics of CMAX.

A. Simulated 4D Planar Pushing in the Presence of Obstacles

In this experiment, the task is for a robotic gripper to push a cube from a start location to a goal location in the presence of static obstacles without any resets, as shown in Figure 2 (right). This can be represented as a planning problem in 4D continuous state space \mathbb{S} with any state represented as the tuple $s = (g_x, g_y, o_x, o_y)$ where (g_x, g_y) are the xy-coordinates of the gripper and (o_x, o_y) are the xy-coordinates of the object. The model \hat{M} used for planning *does not* have the static obstacles and the robot can only discover the state-action pairs that are affected due to the obstacles through real world executions. The action space \mathbb{A} is a discrete set of 4 actions that move the gripper end-effector in the 4 cardinal directions by a fixed offset using an IK-based controller. The cost of each transition is 1 when the object is not at the goal location, and 0 otherwise.

⁶Code to reproduce simulated experiments can be found at <https://github.com/vvanirudh/CMAX>

We compare CMAX with the following baselines: a model-free Q-learning approach [25] that learns from online executions in environment and does not use the model \hat{M} , and a model learning approach that uses limited-expansion search for planning but updates a learned residual that compensates for the discrepancy in dynamics between the model and environment. The model learning approach is very similar to previous works that learn residual dynamics models and have been shown to work well in episodic settings [27, 11, 28]. We chose two function approximators for the learned residual dynamics to account for model learning approaches that use global function approximators such as neural networks (NN) [13], and local function approximators such as K-nearest neighbor regression (KNN) [26, 15]. Finally, we compare against a limited-expansion search planner that uses an accurate model with the full knowledge about obstacles to understand the difficulty of the task. Specific details on the architecture and baseline parameters can be found in Appendix D.

For our implementation, we follow Algorithm 3 with euclidean distance metric, $\xi = 0.01$, and $\delta = 0.02$. These values are chosen to capture the discrepancies observed in the object and gripper position when pushed into an obstacle, and the size of the obstacles. We use the same values for the model learning KNN baseline to ensure a fair comparison. The results of our experiments are presented in Table I. We notice that all the approaches have almost the same performance when both model and environment have no obstacles (first column). This validates that all the baselines do well when the model is accurate. However, when the model is inaccurate (second column), the performance varies across baselines. Q-learning performs decently well since it relies on the model only for the initialized Q-values and not during online executions, but as the task is now more difficult, it solves much fewer trials and is highly suboptimal. It is interesting to see that model learning baselines do not do as well as one would expect. This can be attributed to the number of online executions required to learn the correct residual, which can be prohibitively large. Among the two model learning baselines, KNN works better since it requires fewer samples to learn the residual, while NN requires large amounts of data. In contrast, CMAX does not seek to learn the true dynamics and instead is more focused on reaching the goal quickly. When compared with a planner that uses the accurate model with obstacles and solves 17 trials (last row in Table I), our approach solves 16 trials and achieves the lowest mean number of timesteps to reach the goal among all baselines. We would like to note that the planner with accurate model takes a larger number of timesteps because we used the same initial cost-to-go estimates as other approaches. The initial cost-to-go estimates are more accurate for the model with no obstacles than for the model with obstacles. Hence, it spends a larger number of timesteps updating cost-to-go estimates. This experiment shows that by focusing on reaching the goal and not trying to correct the model dynamics, CMAX performs the best and solves the most number of trials among baselines.

	Steps	% Success
CMAX	47 ± 6	100%
RTAA*	138 ± 65	30%




Fig. 2: (left) Results for simulated 7D arm planning experiment comparing RTAA* and CMAX. Each entry in the Steps column is obtained using 10 trials with random start configurations and goal locations, and we present mean and standard error of number of timesteps it takes the arm to reach the goal *among successful trials*. The % success column indicates percentage of successful trials where the arm reached the goal in less than 300 timesteps. (right) 4D Planar Pushing in the presence of obstacles. The task is to push the black box to the red goal using the end-effector.

B. 3D Pick-and-Place with a Heavy Object

The task of this physical robot experiment (Figure 3) is to pick and place a heavy object using a PR2 arm from a start pick location to a goal place location while avoiding an obstacle. This can be represented as a planning problem in 3D discrete state space \mathbb{S} where each state corresponds to the 3D location of the end-effector. Since it is a relatively small state space, we use exact planning updates without any function approximation following Algorithm 2 with $K = 3$ expansions. The action space is a discrete set of 6 actions corresponding to a fixed offset movement in positive or negative direction along each dimension. The model \hat{M} used by planning *does not* model the object as heavy and hence, does not capture the dynamics of the arm correctly when it holds the heavy object. Specific details regarding the experiment can be found in Appendix E.

We observe that if the object was not heavy, then the arm takes the object from the start pick location to the goal place location on the optimal path which goes above the obstacle (first 3 images of Figure 3). However, when executed with a heavy object, the arm cannot lift the object beyond a certain height as its joint torque limits are reached. At this point, the robot notes the discrepancy in dynamics between the model \hat{M} and the real world, and inflates the cost of any executed transition that tried to move the object higher. Subsequently, the robot figures out an alternate path that does not require it to lift the object higher by taking the object behind the obstacle to the goal place location (last 4 images of Figure 3). The robot takes 36 timesteps (25.8 seconds) to reach the goal with the heavy object, in comparison to 26 timesteps (22.8 seconds) for the light object (see video). Thus, the robot using CMAX successfully completes the task despite having a model with inaccurate dynamics.

C. 7D Arm Planning with a Non-Operational Joint

The task of this physical robot experiment (Figure 4) is to move the PR2 arm with a non-operational joint from a start configuration so that the end-effector reaches a goal location, specified as a 3D region. We represent this as a planning problem in 7D discrete statespace \mathbb{S} where each dimension corresponds to a joint of the arm bounded by its joint limits. The action space \mathbb{A} is a discrete set of size 14 corresponding to moving each joint by a fixed offset in the positive or negative



Fig. 3: Physical robot 3D pick-and-place experiment. The task is to pick the object (light - wooden block, heavy - black dumbbell) and place it at the goal location (green) while avoiding the obstacle (box). For the light object (first 3 images), the model dynamics are accurate and the robot takes it on the optimal path that goes above the obstacle. For the heavy object (next 3 images), the model dynamics are inaccurate, but using CMAX the robot discovers that there is a discrepancy in dynamics when the object is lifted beyond a certain height (due to joint torque limits), adds hyperspheres at that height to account for these transitions (red spheres in the last image), and quickly finds an alternate path going behind the obstacle.




Fig. 4: Physical robot 7D arm planning experiment. The task is to start from a fixed configuration (shown in the first image) and move the arm so that the end-effector reaches the object place location (green). When the shoulder lift joint is operational, the robot uses the joint to quickly find a path to the goal (middle image). However, when the joint is non-operational, it encounters discrepancies in its model and compensates by finding a path that uses other joints to reach the goal (last image.)

direction. The model \hat{M} used for planning *does not* know that a joint is non-operational and assumes that the arm can attain any configuration within the joint limits. In the real world, if the robot tries to move the non-operational joint, the arm does not move. Specific details regarding the experiment can be found in Appendix F.

For the purpose of this experiment since the state space is very large, we follow Algorithm 3 with $\delta = 1$, $\xi = 1$, and make the shoulder lift joint (marked by red cross and arrows in last image of Figure 4) of PR2 non-operational. We use a kernel regressor with RBF kernel of length scale $\gamma = 10$ for the cost-to-go function approximation. Figure 4 shows CMAX operating in the real world to place an object at a desired location with a goal tolerance of 10 cm. When the shoulder lift joint is operational, the robot finds a path quickly to the place location by using the joint (middle image of Figure 4). However, when the shoulder lift joint is non-operational, the robot notes discrepancy in dynamics whenever it tries to move the joint, places hyperspheres in 7D to inflate the cost, and comes up with an alternate path (last image of Figure 4) to reach the place location. The robot takes 13 timesteps (32.4 seconds) to reach the goal location with the non-operational joint, in comparison to 10 timesteps (25.8 seconds) for the case where the joint is working (see video). Thus, the robot successfully finds a path to the place location despite using a model with inaccurate dynamics.

To emphasize the need for cost-to-go function approximation and local generalization from hyperspheres in large state spaces, we compared CMAX against RTAA*, an exact planning method that uses a tabular representation for cost-to-go estimates and updates model dynamics online. Results are presented in Figure 2 (left) and show that RTAA* fails to solve 7 of the 10 trials whereas CMAX solves all of them, and

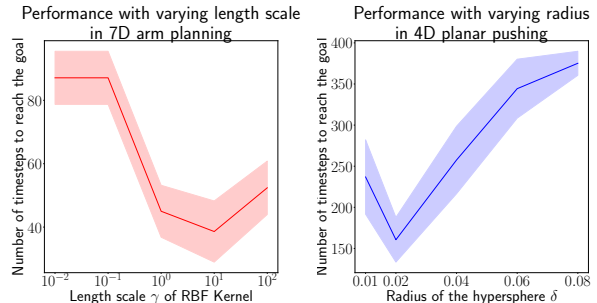


Fig. 5: (left) Performance of CMAX for 7D arm planning as the smoothness of the cost-to-go function approximator varies. The plot is generated for each value of length scale γ by generating 10 random start configurations and goal locations, and running our approach for a maximum of 100 timesteps. (right) Performance of our approach for 4D planar pushing as the radius of the hypersphere δ varies. The plot is generated for each value of radius δ by generating 10 random start and goal locations, and running CMAX for a maximum of 400 timesteps.

in smaller mean number of timesteps.

D. Effect of Function Approximation and Size of Hyperspheres

While previous experiments have tested CMAX against other baselines and on a physical robot, this experiment is designed to evaluate the effect of cost-to-go function approximation and the size of hyperspheres on the performance of CMAX in large state spaces (Algorithm 3.) For the first set of experiments (Figure 5 left), we use the setup of Section V-C and focus on varying the smoothness of the kernel regressor cost-to-go function approximation by varying the length scale γ of the RBF kernel. Intuitively, small length scales result in approximation with high variance, and for large scales we obtain highly smooth approximation. We notice that for small γ , the performance is poor and as γ increases, the performance of CMAX becomes better as it can generalize the cost-to-go estimates in the state space. However, for large γ the performance deteriorates as it fails to capture the difference in cost-to-go values among nearby states due to excessive smoothing. This showcases the need for generalization in cost-to-go estimates for efficient updates in large state spaces.

For the second set of experiments (Figure 5 right), we vary the radius of the hyperspheres δ introduced whenever an incorrect state-action pair is discovered in Algorithm 3. We use the setup of Section V-A, vary δ and observe the number of timesteps it takes the robot to push the object to the goal. We observe that when δ is large, the performance is poor as we potentially penalize state-action pairs that are not incorrect and could result in a very suboptimal path. However,

% Ice	0%	40%	80%
CMAX	78 ± 4	231 ± 18	2869 ± 331
RTAA*	78 ± 4	219 ± 18	2185 ± 249
Q-Learning	3914 ± 303	1220 ± 103	996 ± 108

TABLE II: Results for gridworld navigation in presence of icy states for a grid of size 100×100 . Each entry is obtained using 50 random seeds, and we present the mean and standard error of the number of timesteps it takes the robot to reach the goal. The columns represent the percentage of icy states in the gridworld.

a very small δ can also lead to a poor performance, as we need more online executions to discover the set of incorrect state-action pairs. Hence, the radius δ needs to be chosen carefully to quickly “cover” the incorrect set, while not penalizing any correct state-action pairs.

E. Simulated 2D Gridworld Navigation with Icy States

In our final experiment, we want to understand the performance of CMAX compared to other baselines in small domains where model dynamics can be represented using a table, and can be updated efficiently. We consider the 2D gridworld such as the one shown in Figure 1(right) with icy states where the robot slips (moving left or right on ice moves the robot by two cells.) The model used for planning *does not* contain ice, and is an empty gridworld. The results are presented in Table II. We can observe that model-free approaches like Q-learning perform well compared to model-based approaches in cases where the model available is highly inaccurate (see Table II last column.) However, when the model is reasonably accurate RTAA* performs the best. But the results show that even in domains where model dynamics are simple and can be updated efficiently, CMAX competes closely with RTAA*. Thus, our approach is still applicable in such domains and is relatively easier to implement.

VI. RELATED WORK

The proposed approach has components concerning real-time heuristic search, local function approximation methods, and dealing with inaccuracy in models. There is a wide array of existing work at the intersection of planning and learning that deal with these topics. Notably, we leverage prior work on real-time heuristic search [22, 20] for the limited-expansion search-based planner presented in Algorithm 1. Using local function approximation methods in robotics has been heavily explored in seminal works [32, 2] due to their smaller sample complexity requirements and local generalization properties that do not cause interference [3, 8]. More recently, [15], [26] and [5] have also proposed approaches that learn local models from online executions. However unlike CMAX, they use these models to approximate the dynamics of the real world. Our work is also closely related to the field of real-time reinforcement learning that tackles the problem of acting near-optimally in unknown environments, without any resets [29, 4, 21]. The analysis presented in Theorem 4.1 and 4.2 borrows several useful results from Koenig and Simmons [21]. Prior works in model-based reinforcement learning with provable guarantees, such as [17, 5, 7, 16], are also related. However, these works learn the true dynamics by updating

the model and give sample complexity results in the finite-horizon setting or discounted infinite-horizon setting, unlike our shortest path setting. Among these works, Kakade et al. [16], which proposes a method for exploration in metric state spaces, serves as an inspiration for the covering number bounds given in Theorem 4.2. The work that is most closely related to ours is Jiang [14] which proposed an approach that uses a similar idea of updating the cost function, in cases where updating the model dynamics is infeasible. However, their approach is suitable only for episodic settings and small state spaces. Concurrent work by McConachie et al. [24] employs a binary classifier trained on offline data to predict whether a transition is incorrect or correct, that is then queried during online motion planning to construct the search tree consisting of only transitions that are classified as correct.

VII. DISCUSSION AND CONCLUSION

CMAX is the first approach for interleaving planning and execution that does not require updating dynamics, and is guaranteed to reach the goal despite using an inaccurate dynamical model. The biggest advantage of CMAX is that it does not rely on any knowledge of how the model is inaccurate, and whether it can be updated in real-time. Hence, it is broadly applicable in real world robotic tasks with complex inaccurate models. In domains where modeling the true dynamics is intractable, such as deformable manipulation, CMAX can still be employed to ensure successful execution. In comparison, approaches that update the model dynamics online rely on the flexibility of the model to be updated, knowledge of what is lacking in the model, and a large number of online executions to correct it. For example, to learn accurate dynamics for a transition in N -D statespace we need at least N samples in the worst case, whereas our approach needs only 1 sample to observe a discrepancy and inflate the cost. The most important shortcomings of CMAX are Assumptions 4.1 and 4.2, which are hard to verify, and are not satisfied in several real world robotic tasks (more details in Appendix C). For example, consider the task of opening a spring-loaded door which is not modeled as loaded. All transitions would have discrepancy in dynamics, and CMAX as is would fail at completing the task in a reasonable amount of time. In addition, the hyperparameter δ describing the radius of hypersphere needs to be tuned carefully for each domain which is a limitation of CMAX.

To summarize, we present CMAX for interleaving planning and execution using inaccurate models that does not require updating the dynamics of the model, and still provably completes the task. We propose practical algorithms for both small and large state spaces, and deploy them successfully in real world robot tasks showing its broad applicability. In simulation, we analyze CMAX and show that it outperforms baselines that update dynamics online. Future directions include establishing similar guarantees like Theorem 4.2 in the approximate planning setting, and relaxing Assumptions 4.1, 4.2 so that CMAX is applicable to a wider range of robotic tasks.

REFERENCES

- [1] Pieter Abbeel, Morgan Quigley, and Andrew Y. Ng. [Using inaccurate models in reinforcement learning](#). In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, pages 1–8, 2006. doi: 10.1145/1143844.1143845.
- [2] Christopher G. Atkeson and Stefan Schaal. [Learning tasks from a single demonstration](#). In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation, Albuquerque, New Mexico, USA, April 20-25, 1997*, pages 1706–1712, 1997. doi: 10.1109/ROBOT.1997.614389.
- [3] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. [Locally Weighted Learning for Control](#). *Artif. Intell. Rev.*, 11(1-5):75–113, 1997. doi: 10.1023/A:1006511328852.
- [4] Andrew G. Barto, Steven J. Bradtke, and Satinder P. Singh. [Learning to Act Using Real-Time Dynamic Programming](#). *Artif. Intell.*, 72(1-2):81–138, 1995. doi: 10.1016/0004-3702(94)00011-O.
- [5] Andrey Bernstein and Nahum Shimkin. [Adaptive-resolution reinforcement learning with polynomial exploration in deterministic domains](#). *Machine Learning*, 81(3):359–397, 2010. doi: 10.1007/s10994-010-5186-7.
- [6] Dimitri P. Bertsekas. *Dynamic programming and optimal control, 3rd Edition*. Athena Scientific, 2005. ISBN 1886529264.
- [7] Ronen I. Brafman and Moshe Tennenholtz. [R-MAX - A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning](#). *J. Mach. Learn. Res.*, 3:213–231, 2002.
- [8] Adam Coates, Pieter Abbeel, and Andrew Y. Ng. [Learning for control from multiple demonstrations](#). In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008*, pages 144–151, 2008. doi: 10.1145/1390156.1390175.
- [9] Benjamin J. Cohen, Gokul Subramania, Sachin Chitta, and Maxim Likhachev. [Planning for Manipulation with Adaptive Motion Primitives](#). In *IEEE International Conference on Robotics and Automation, ICRA 2011, Shanghai, China, 9-13 May 2011*, pages 5478–5485, 2011. doi: 10.1109/ICRA.2011.5980550.
- [10] Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. [Gaussian Processes for Data-Efficient Learning in Robotics and Control](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(2):408–423, 2015. doi: 10.1109/TPAMI.2013.218.
- [11] Sehoon Ha and Katsu Yamane. [Reducing hardware experiments for model learning and policy optimization](#). In *IEEE International Conference on Robotics and Automation, ICRA 2015, Seattle, WA, USA, 26-30 May, 2015*, pages 2620–2626, 2015. doi: 10.1109/ICRA.2015.7139552.
- [12] Kris K. Hauser, Timothy Bretl, Kensuke Harada, and Jean-Claude Latombe. [Using Motion Primitives in Probabilistic Sample-Based Planning for Humanoid Robots](#). In *Algorithmic Foundation of Robotics VII, Selected Contributions of the Seventh International Workshop on the Algorithmic Foundations of Robotics, WAFR 2006, July 16-18, 2006, New York, NY, USA*, pages 507–522, 2006. doi: 10.1007/978-3-540-68405-3_32.
- [13] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. [When to Trust Your Model: Model-Based Policy Optimization](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 12498–12509, 2019.
- [14] Nan Jiang. [PAC Reinforcement Learning With an Imperfect Model](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3334–3341, 2018.
- [15] Nicholas K. Jong and Peter Stone. [Model-based function approximation in reinforcement learning](#). In *6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007), Honolulu, Hawaii, USA, May 14-18, 2007*, page 95, 2007. doi: 10.1145/1329125.1329242.
- [16] Sham M. Kakade, Michael J. Kearns, and John Langford. [Exploration in Metric State Spaces](#). In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pages 306–312, 2003.
- [17] Michael J. Kearns and Satinder P. Singh. [Near-Optimal Reinforcement Learning in Polynomial Time](#). *Machine Learning*, 49(2-3):209–232, 2002. doi: 10.1023/A:1017984413808.
- [18] Michael J. Kearns, Yishay Mansour, and Andrew Y. Ng. [A Sparse Sampling Algorithm for Near-Optimal Planning in Large Markov Decision Processes](#). *Mach. Learn.*, 49(2-3):193–208, 2002. doi: 10.1023/A:1017932429737.
- [19] Diederik P. Kingma and Jimmy Ba. [Adam: A Method for Stochastic Optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [20] Sven Koenig and Maxim Likhachev. [Real-time adaptive A*](#). In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*, pages 281–288, 2006. doi: 10.1145/1160633.1160682.
- [21] Sven Koenig and Reid G. Simmons. [Complexity Analysis of Real-Time Reinforcement Learning](#). In *Proceedings of the 11th National Conference on Artificial Intelligence. Washington, DC, USA, July 11-15, 1993*,

- pages 99–107, 1993.
- [22] Richard E. Korf. [Real-Time Heuristic Search](#). *Artif. Intell.*, 42(2-3):189–211, 1990. doi: 10.1016/0004-3702(90)90054-4.
 - [23] Steven M. LaValle and James J. Kuffner Jr. [Randomized Kinodynamic Planning](#). *I. J. Robotics Res.*, 20(5):378–400, 2001. doi: 10.1177/02783640122067453.
 - [24] Dale McConachie, Thomas Power, Peter Mitrano, and Dmitry Berenson. [Learning When to Trust a Dynamics Model for Planning in Reduced State Spaces](#). *IEEE Robotics Autom. Lett.*, 5(2):3540–3547, 2020. doi: 10.1109/LRA.2020.2972858.
 - [25] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumar, Daan Wierstra, Shane Legg, and Demis Hassabis. [Human-level control through deep reinforcement learning](#). *Nature*, 518(7540):529–533, 2015. doi: 10.1038/nature14236.
 - [26] Ali Nouri and Michael L. Littman. [Multi-resolution Exploration in Continuous Spaces](#). In *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 8-11, 2008*, pages 1209–1216, 2008.
 - [27] Divyam Rastogi, Ivan Koryakovskiy, and Jens Kober. [Sample-efficient reinforcement learning via difference models](#). In *Machine Learning in Planning and Control of Robot Motion Workshop at ICRA*, 2018.
 - [28] Matteo Saveriano, Yuchao Yin, Pietro Falco, and Dongheui Lee. [Data-efficient control policy search using residual dynamics learning](#). In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 4709–4715, 2017. doi: 10.1109/IROS.2017.8206343.
 - [29] Richard S. Sutton. [Dyna, an Integrated Architecture for Learning, Planning, and Reacting](#). *SIGART Bulletin*, 2(4):160–163, 1991. doi: 10.1145/122344.122377.
 - [30] Emanuel Todorov, Tom Erez, and Yuval Tassa. [MuJoCo: A physics engine for model-based control](#). In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012, Vilamoura, Algarve, Portugal, October 7-12, 2012*, pages 5026–5033, 2012. doi: 10.1109/IROS.2012.6386109.
 - [31] Anirudh Vemula, Katharina Mülling, and Jean Oh. [Modeling cooperative navigation in dense human crowds](#). In *2017 IEEE International Conference on Robotics and Automation, ICRA 2017, Singapore, Singapore, May 29 - June 3, 2017*, pages 1685–1692, 2017. doi: 10.1109/ICRA.2017.7989199.
 - [32] Sethu Vijayakumar and Stefan Schaal. [Locally Weighted Projection Regression: Incremental Real Time Learning in High Dimensional Space](#). In *Proceedings of the Sev-*

enteenth International Conference on Machine Learning (ICML 2000), Stanford University, Stanford, CA, USA, June 29 - July 2, 2000, pages 1079–1086, 2000.