

Lazy Heuristic Search for Solving POMDPs with Expensive-to-Compute Belief Transitions

Muhammad Suhail Saleem, Rishi Veerapaneni, Maxim Likhachev

Carnegie Mellon University
{msaleem2, rveerapa, mlikhach}@andrew.cmu.edu

Abstract

Heuristic search solvers like RTDP-Bel and LAO* have proven effective for computing optimal and bounded sub-optimal solutions for Partially Observable Markov Decision Processes (POMDPs), which are typically formulated as belief MDPs. A belief represents a probability distribution over possible system states. Given a parent belief and an action, computing belief state transitions involves Bayesian updates that combine the transition and observation models of the POMDP to determine successor beliefs and their transition probabilities. However, there is a class of problems, specifically in robotics, where computing these transitions can be prohibitively expensive due to costly physics simulations, raycasting, or expensive collision checks required by the underlying transition and observation models, leading to long planning times. To address this challenge, we propose Lazy RTDP-Bel and Lazy LAO*, which defer computing expensive belief state transitions by leveraging Q -value estimation, significantly reducing planning time. We demonstrate the superior performance of the proposed lazy planners in domains such as contact-rich manipulation for pose estimation, outdoor navigation in rough terrain, and indoor navigation with a 1-D LiDAR sensor. Additionally, we discuss practical Q -value estimation techniques for commonly encountered problem classes that our lazy planners can leverage. Our results show that lazy heuristic search methods dramatically improve planning speed by postponing expensive belief transition evaluations while maintaining solution quality.

Extended version — <https://arxiv.org/abs/2506.00285>

1 Introduction

Partially Observable Markov Decision Processes (POMDPs) provide a principled framework for decision-making under uncertainty and have been widely used in robotics, including autonomous navigation and manipulation tasks (Brechtel, Gindele, and Dillmann 2014; Pajarinen and Kyrki 2017). As the true state of the system is not directly observable, POMDPs maintain a probability distribution over possible system states, referred to as belief states. They are commonly formulated and solved as a Markov Decision Process (MDP) in the belief space. However, solving POMDPs is notoriously challenging due to the belief space growing exponentially with each added state variable (curse of dimension-

ality), as well as with the planning horizon (curse of history) (Kaelbling, Littman, and Cassandra 1998).

Search-based solvers like RTDP-Bel (Bonet and Geffner 2009) and LAO* (Hansen and Zilberstein 2001) leverage heuristics to focus computational efforts on promising regions of the belief space, efficiently converging to optimal or bounded suboptimal solutions. Although powerful and effective in many problems, their applicability still remains limited—particularly in real-world robotics problems. One of the many reasons for this is the inherent assumption that belief transitions are readily available. Computing transitions in the belief space, i.e., determining the successor beliefs and their probabilities given the current belief over the system state and an action, involves Bayesian updates that combine the transition and observation models of the underlying POMDP. However, in many robotics applications, this computation poses a significant bottleneck due to the high computational cost of querying these models. For instance, outdoor robot navigation may require high-fidelity physics simulations to predict action outcomes, while contact-rich manipulation tasks often involve expensive collision checks to model expected interactions. These operations are expensive to query, even for a single system state. Since the belief state is a distribution over the system states, computing the transitions from a belief state would involve querying these operations for *every* state in (the support of) the belief. Since in a typical problem we can expect to search over thousands of belief states, the total cost of computing belief transitions becomes prohibitively expensive, limiting the practicality of POMDP-based solutions for such applications.

To mitigate this challenge, we propose lazy heuristic search solvers for POMDPs, which defer expensive belief state transition computations until necessary. Specifically, we introduce two novel algorithms: Lazy RTDP-Bel and Lazy LAO*. These methods leverage Q -value estimation to identify the most promising action from a belief state, and only the transition corresponding to the most promising action is computed. The belief state transitions for other actions are postponed until the search deems these actions necessary, i.e., have the potential to improve the current solution. By postponing these expensive evaluations and only computing them when needed, we save significant computation time. We show that the Q -value estimates leveraged by our lazy planners are equivalent to heuristic functions and are readily available for common problems, particularly in

robotics. Moreover, if the Q -value estimates are conservative, then the proposed lazy planners, like their vanilla counterparts, are guaranteed to converge to the optimal solution.

We demonstrate the effectiveness of our proposed lazy planners on three robotics problems i) Manipulation for pose estimation using contacts, ii) Indoor navigation with a 1D LiDAR sensor, and iii) Outdoor rough terrain navigation. Belief transition computations for each of these domains are expensive due to them requiring mesh-to-mesh collision checks, ray casting operations, and physics-based simulations, respectively. Results presented in Section 5, demonstrate the superior performance of the lazy algorithms on the abovementioned problems, significantly improving planning efficiency while maintaining solution quality.

2 Background

In this manuscript, we introduce the concept of laziness in the context of Goal-POMDPs (Bertsekas 2012), though it naturally extends to reward-based formulations.

2.1 Goal-POMDPs

Formally, a discrete Goal-POMDP is defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{Z}, G, \mathcal{T}, \mathcal{O}, c, b_0 \rangle$, where \mathcal{S} , \mathcal{A} , and \mathcal{Z} are finite non-empty sets representing the state space, action space, and observation space respectively. $G \subseteq \mathcal{S}$ is the set of goal states. The transition model $\mathcal{T}(s, a, s') = p(s'|s, a)$ defines the probability of transitioning to state s' when executing action a from state s , while the observation model $\mathcal{O}(s', a, z) = p(z|s', a)$ gives the probability of observing z upon entering s' via action a . The cost function $c(s, a)$ determines the cost of executing action a from state s , and b_0 is the initial belief state, i.e., the distribution over possible initial states. We assume that goal states $g \in G$ are cost-free, absorbing, and fully observable ($z_g \in \mathcal{Z}$ denotes the observation made from a goal state): (i) $c(g, a) = 0, \forall a \in \mathcal{A}$, (ii) $\mathcal{T}(g, a, g) = 1$, and (iii) $\mathcal{O}(g, a, z_g) = 1$ for all $a \in \mathcal{A}$, with $\mathcal{O}(s, a, z_g) = 0$ for $s \notin G$.

POMDPs are commonly solved by reformulating them as completely observable MDPs over the belief space. A belief state b is a probability distribution over states, where $b(s)$ represents the probability of being in state s . Although the system state is not directly observable, we can compute belief updates that capture the effects of actions and observations on the belief state. The belief update after executing action a from belief b is given by

$$b_a(s) = P(s|b, a) = \sum_{s' \in \mathcal{S}} \mathcal{T}(s', a, s) b(s') \quad (1)$$

Similarly, the probability of observing z after taking action a from b is,

$$b_a(z) = P(z|b, a) = \sum_{s \in \mathcal{S}} b_a(s) \mathcal{O}(s, a, z) \quad (2)$$

Then, we can write the belief update when executing a from b and observing z as

$$b_a^z(s) = \mathcal{O}(s, a, z) b_a(s) / b_a(z) \quad \text{if } b_a(z) \neq 0 \quad (3)$$

This formulation transforms the POMDP into a fully observable belief MDP, where the transition function is given

Algorithm 1: RTDP-BEL

```

1: while NOT CONVERGED do
2:    $b = b_0$ 
3:   Sample state  $s$  with probability  $b(s)$ 
4:   while  $b$  is not a goal belief do
5:     Evaluate each action  $a \in \mathcal{A}$  from  $b$ 
6:     Compute value of each action as:
       
$$Q(b, a) = c(b, a) + \sum_{z \in \mathcal{Z}} P(z|b, a) V(b_a^z)$$

       
$$\triangleright V(b_a^z) = \text{heur}(b_a^z) \text{ if uninitialized}$$

7:     Select action  $a_{best}$  that minimizes  $Q(b, a)$ 
8:     Update value  $V(b) = Q(b, a_{best})$ 
9:     Sample  $s'$  with probability  $\mathcal{T}(s, a_{best}, s')$ 
10:    Sample  $z$  with probability  $\mathcal{O}(s', a_{best}, z)$ 
11:    Compute  $b_a^z$  and set  $b := b_a^z$  and  $s := s'$ 

```

by $P(b_a^z|b, a) = b_a(z)$. Consequently, reaching a goal state in the original POMDP translates to reaching a goal belief b_g , where $b_g(s) = 0$ for all $s \notin G$. The objective of the problem is to compute a belief state policy from the initial belief b_0 to a goal belief b_g while minimizing the expected cost. The Bellman equation for the belief MDP is given by

$$V^*(b) = \min_{a \in \mathcal{A}} \left\{ c(b, a) + \sum_{z \in \mathcal{Z}} b_a(z) V^*(b_a^z) \right\} \quad (4)$$

for non-goal beliefs and $V^*(b_g) = 0$ for goal beliefs. $c(b, a)$ is the expected cost of executing action a from b , i.e., $c(b, a) = \sum_{s \in \mathcal{S}} c(s, a) b(s)$.

2.2 RTDP-Bel

RTDP-Bel¹ (Bonet and Geffner 2009), outlined in Alg 1, is an adaptation of RTDP (Barto, Bradtke, and Singh 1995) to POMDPs. As an asynchronous value iteration algorithm, it converges to the optimal value function and policy for only the relevant regions of the belief space. The algorithm operates through a series of rollouts (or trials), each beginning from the initial belief state b_0 and continuing until a goal belief is reached. At each belief state b , all possible actions are evaluated. Evaluating an action a from b involves computing the belief transitions corresponding to the action, i.e., the successor beliefs b_a^z . The utility of these actions is then assessed by computing their Q -values (Line 6). The action with the lowest Q -value is selected, and the value estimate of b is updated accordingly (Lines 7–8). The process then proceeds to a successor belief, determined based on the chosen action and a randomly sampled observation (Lines 9–11).

A crucial component of the algorithm's performance is the heuristic function, heur , which initializes the value estimates. A well-informed heuristic accelerates convergence by guiding the search toward the goal quickly. The heuristic must be admissible, i.e., underestimate the optimal value function, for the converged policy to be optimal. Further, Kim, Salzman, and Likhachev (2019) show that inflating an

¹We discuss the unapproximated version of RTDP-Bel that maintains the values for each belief state independently.

Algorithm 2: LAO*

```
1:  $G^* \leftarrow \{b_0\}$ 
2: while  $G^*$  has non-terminal tip states do
3:   Identify a non-terminal tip state  $b$  in  $G^*$ 
4:   Evaluate each action  $a \in \mathcal{A}$  from  $b$ 
5:   Compute value of each action as:
      
$$Q(b, a) = c(b, a) + \sum_{z \in \mathcal{Z}} P(z|b, a)V(b_a^z)$$

      
$$\triangleright V(b_a^z) = \text{heur}(b_a^z) \text{ if uninitialized}$$

6:   Select action  $a_{best}$  that minimizes  $Q(b, a)$ 
7:   Update value  $V(b) = Q(b, a_{best})$ 
8:   Create set  $Z$  with  $b$  and ancestors of  $b$  in  $G^*$ 
9:   IMPROVEVALUES( $Z$ )  $\triangleright$  Value iteration
10:  Update  $G^*$  to current best solution
11: IMPROVEVALUES( $G^*$ )  $\triangleright$  Value iteration
12: if  $G^* \neq$  current best solution then
13:   Update  $G^*$  to current best solution
14:   Goto Line 2
15: return  $G^*$ 
```

admissible heuristic by ϵ allows each search iteration to be more strongly influenced by the heuristic, leading to significantly faster convergence to an ϵ -suboptimal solution.

2.3 LAO*

LAO* (Hansen and Zilberstein 2001) outlined in Alg 2, interleaves forward expansion with dynamic backup updates while ensuring that no states unreachable from the start belief b_0 under the current best policy are considered. It maintains a partial solution graph G^* , which consists of all belief states reachable from b_0 under the current policy. The algorithm iteratively expands and updates G^* until no non-terminal tip states remain.

A belief state b is classified as a non-terminal tip state if it is neither a goal belief nor has been expanded, meaning its successor beliefs b_a^z have not been generated. When such a state is encountered, all possible actions from b are evaluated, and their corresponding successor beliefs are computed. The optimal action is selected based on the Q -value, and the value of b is updated accordingly (Lines 4–7). Next, a set Z is formed, including b and its ancestor beliefs in G^* , and value iteration is performed on Z to update both values and optimal actions (Line 9). Finally, the solution graph G^* is updated (Line 10). This process repeats until no non-terminal tip states remain in G^* .

Once no further updates are required, value iteration is performed on all belief states in G^* . If the updates lead to changes in the best solution, G^* is revised, and the algorithm returns to Line 2². If the values converge, the algorithm terminates, returning the final solution policy represented by G^* (Lines 11-15). Similar to RTDP-Bel, the heuristic guides the search toward the most promising belief regions, facilitating faster convergence. The same principles of optimal-

²This is a slightly modified interpretation of LAO* to facilitate explanation, though it maintains the same fundamental properties as the original algorithm.

ity with an admissible heuristic and bounded sub-optimality with an inflated heuristic apply to LAO* as well.

3 Related Work

Beyond heuristic search, two widely studied approaches for scalably solving POMDPs are point-based methods and Monte Carlo methods. Point-based POMDP solvers are a class of approximate algorithms designed to efficiently solve large POMDPs by selectively maintaining and updating a representative set of belief points. Solvers such as PBVI (Pineau et al. 2003), HSVI (Smith and Simmons 2012), and SARSOP (Kurniawati, Hsu, and Lee 2008), leverage value iteration on a subset of beliefs to compute near-optimal policies while significantly reducing computational complexity. However, Point-based backups—a core operation of these solvers—assume the transition and observation models are readily accessible and require querying them at every state in the state space. While feasible for simpler problems, where one can precompute a tabulated transition and observation matrix, our domains involve large state spaces with expensive models, making such enumerations impractical. Extensive discussions on the many point-based solvers can be found in Shani, Pineau, and Kaplow (2013). On the other hand, Monte Carlo methods, particularly suited for online planning, approximate the value of belief states by simulating action sequences and propagating rewards. Works like Silver and Veness (2010); Lee et al. (2018); Kurniawati and Yadav (2016) sample state trajectories and update action values based on rollouts. DESPOT (Somani et al. 2013), a popular approach in this realm, improves efficiency by maintaining a fixed number of sampled scenarios, generating a sparsely sampled belief tree that performs well in problems with large observation spaces.

In deterministic search, the concept of laziness has been widely explored to defer checking the validity of edges, which is often the computational bottleneck, especially in robotics applications where collision checking is expensive. Lazy Weighted A* (LWA*) (Cohen, Phillips, and Likhachev 2015) postpones edge evaluations until expansion, meaning only edges leading to expanded nodes are evaluated. On the other hand, LazySP (Dellin and Srinivasa 2016) initially assumes all edges are valid and finds a path to the goal; if an edge is found to be invalid, it is removed from the graph, and the search is rerun, minimizing the number of evaluations. Variants presented in Mandalika, Salzman, and Srinivasa (2018); Bhardwaj et al. (2021); Mandalika et al. (2019) further leverage laziness to accelerate planning. While prior work has focused on deferring edge evaluations in deterministic settings, we extend this concept to stochastic domains by postponing expensive belief transition computations.

4 Methodology

Many heuristic search methods for POMDPs are on-policy, exploring only belief states in the current best solution policy (Bonet and Geffner 2009; Hansen and Zilberstein 2001; Ross, Chaib-Draa et al. 2007). They maintain Q -values for actions available from each encountered belief state. At a given belief b , the best action a_{best} , minimizing the Q -values is selected, and the search explores only its successor beliefs.

The Q -value of a_{best} is then updated based on the outcome of this exploration. If this update changes the minimizing action at b , a_{best} is updated, and the search proceeds with the new best action and its corresponding successor beliefs.

The key observation here is that, since only the successor beliefs of the current best action a_{best} are relevant and explored, belief transitions for all other actions remain unused. Consequently, computing these unused transitions introduces unnecessary computational overhead. Hence, we propose to only evaluate the best action a_{best} from any explored belief state. By generating belief transitions only for a_{best} , we defer computing the belief transitions for all other actions until the search deems these actions to be promising, i.e., has potential to be part of the solution policy. This ensures that belief transition evaluations are performed only when necessary, significantly reducing computational complexity. Actions that are never considered promising at any point during the search remain unevaluated. This key principle of postponing action evaluations until they are deemed useful forms the essence of lazy search.

However, existing methods typically initialize Q -values using a one-step lookahead (Alg 1 Line 6; Alg 2 Line 5), which involves evaluating each action by computing belief transitions and using successor belief values (initialized with a heuristic) to determine the actions' Q -values. This process incurs expensive belief transition computations for all actions, which the lazy search aims to defer. Hence, we propose replacing the one-step lookahead with a fast-to-query estimator that does not require belief transition computations. We denote this estimator as \hat{Q}_{init} , as it is used solely for initializing Q -values. Under this approach, only the belief transitions for a_{best} are computed.

The Q -estimators our lazy planners leverage are similar to heuristic functions and can be easily constructed for common problems. Just as search algorithms like RTDP-Bel and LAO* converge optimally with an admissible heuristic, $\text{heur}(b) \leq V^*(b)$, their lazy counterparts guarantee optimality if \hat{Q}_{init} is admissible, i.e., $\hat{Q}_{init}(b, a) \leq Q^*(b, a)$.

We first introduce the lazy adaptations of RTDP-Bel and LAO*, followed by a brief discussion of simple techniques for constructing effective estimators.

4.1 Lazy RTDP-Bel

The lazy adaptation of RTDP-Bel is outlined in Algorithm 3 and is fairly straightforward with key differences being from Lines 7-12. When a belief state is encountered for the first time, the Q -values of all available actions are initialized using the fast-to-compute estimator, \hat{Q}_{init} (Line 7). The algorithm then identifies the best action, i.e., the action that minimizes the Q -values. Unlike standard RTDP-Bel, only this best action is evaluated, meaning that the belief transitions corresponding to the selected action are computed. The Q -value of the evaluated action is then updated accordingly through routine UPDATE Q-VALUES. If this update results in a change in the best action, a_{best} is updated, and the process is repeated until the best action remains unchanged following an update (Lines 8–12). This approach ensures that the algorithm evaluates only the actions that minimize the Q -value for any belief state encountered. By deferring the

Algorithm 3: LAZY RTDP-BEL

```

1: procedure LAZY RTDP-BEL
2:   while NOT CONVERGED do
3:      $b = b_0$ 
4:     Sample state  $s$  with probability  $b(s)$ 
5:     while  $b$  is not a goal belief do
6:       if  $Q$ -values not initialized for  $b$  then
7:         Set  $Q(b, a) = \hat{Q}_{init}(b, a) \forall a \in \mathcal{A}$ 
8:       do  $\triangleright$  Do while loop (e.g., to Line 12)
9:         Select  $a_{best}$  that minimizes  $Q(b, a)$ 
10:        Evaluate  $a_{best}$  from  $b$  if unevaluated
11:        UPDATE Q-VALUES( $b$ )
12:        while  $a_{best} \neq \arg \min_{a \in \mathcal{A}} Q(b, a)$ 
13:          Update value  $V(b) = Q(b, a_{best})$ 
14:          Sample  $s'$  with probability  $\mathcal{T}(s, a_{best}, s')$ 
15:          Sample  $z$  with probability  $\mathcal{O}(s', a_{best}, z)$ 
16:          Compute  $b_{a_{best}}^z$ ; set  $b := b_{a_{best}}^z$  and  $s := s'$ 
17: procedure UPDATE Q-VALUES( $b$ )
18:   for  $a \in \mathcal{A}$  do
19:     if  $a$  evaluated from  $b$  then
20:        $Q(b, a) = \mathcal{C}(b, a) + \sum_{z \in \mathcal{Z}} P(z|b, a)V(b_a^z)$ 
21:      $\triangleright V(b_a^z) = \text{heur}(b_a^z)$  if uninitialized

```

evaluation of actions until they are deemed promising and potentially part of the (bounded sub-)optimal solution, the method efficiently reduces unnecessary computations.

4.2 Lazy LAO*

The lazy adaptation of LAO* outlined in Alg 4 follows the same fundamental structure as vanilla LAO*. Like its standard counterpart, the algorithm maintains the current best partial solution graph, G^* , and iteratively identifies non-terminal tip states for expansion. However, we slightly modify the definition of a non-terminal tip state. In vanilla LAO*, a non-terminal tip state is a non-goal belief whose actions have not been evaluated. In our adaptation, a non-terminal tip state is defined as either (i) a non-goal belief with no evaluated actions or (ii) a non-goal belief whose best action—the action with the lowest Q -value—has not yet been evaluated.

Using this modified definition, we identify a non-terminal tip state in G^* (Line 4). If the belief state has not been encountered before, its Q -values are initialized using the quick-to-compute estimator \hat{Q}_{init} (Line 6). As in Lazy RTDP-Bel, the algorithm identifies the action minimizing the Q -values and evaluates only this action. The evaluated action's value is then updated accurately. If this update changes the best action, the newly identified best action is evaluated, and the process repeats (Lines 7-11). This ensures only promising actions from a state are evaluated. After this step, the algorithm proceeds like LAO*, using the Q -value of the best action to update b . The set Z , consisting of b and its ancestors within G^* , is then constructed, and value improvements are performed. This continues until no non-terminal tip states remain in G^* . At this point, the values of all states in G^* are further refined until either G^* is modified or the values converge (Lines 12-20).

Algorithm 4: Lazy LAO*

```

1: procedure LAZY LAO*
2:    $G^* \leftarrow \{b_0\}$ 
3:   while  $G^*$  has non-terminal tip states do
4:     Identify a non-terminal tip state  $b$  in  $G^*$ 
5:     if  $Q$ -values not initialized for  $b$  then
6:       Set  $Q(b, a) = \hat{Q}_{init}(b, a) \quad \forall a \in \mathcal{A}$ 
7:       do  $\triangleright$  Do while loop (e.g., to Line 11)
8:         Select action  $a_{best}$  that minimizes  $Q(b, a)$ 
9:         Evaluate  $a_{best}$  from  $b$  if unevaluated
10:        UPDATE  $Q$ -VALUES( $b$ )
11:        while  $a_{best} \neq \arg \min_{a \in \mathcal{A}} Q(b, a)$ 
12:          Update value  $V(b) = Q(b, a_{best})$ 
13:          Create set  $Z$  with  $b$  and ancestors of  $b$  in  $G^*$ 
14:          IMPROVEVALUES( $Z$ )
15:          Update  $G^*$  to current best solution
16:        IMPROVEVALUES( $G^*$ )
17:        if  $G^* \neq$  current best solution then
18:          Update  $G^*$  to current best solution
19:          Goto Line 3
20:        return  $G^*$ 

21: procedure IMPROVEVALUES( $\mathcal{B}$ )
22:   while Error bound greater than  $\epsilon$  do
23:     for  $b \in \mathcal{B}$  do
24:       UPDATE  $Q$ -VALUES( $b$ )
25:       Select action  $a_{best}$  that minimizes  $Q(b, a)$ 
26:       if  $a_{best}$  not evaluated from  $b$  then
27:         return
28:       Update value  $V(b) = Q(b, a_{best})$ 

```

A key difference from vanilla LAO* lies in how values are improved in the **IMPROVEVALUES** routine. In standard LAO*, this routine is simple value iteration, assuming all actions from a belief state are either evaluated or can be cheaply evaluated. In contrast, our adaptation accounts for lazy evaluations. If, during backup, the best action of a state changes to an unevaluated one, backups terminate (Line 26). If the best action for any state in either set Z or G^* changes, it implies that the current best partial solution has been modified, necessitating an update to G^* (Lines 15, 18), followed by which the algorithm returns to Line 3 (Lines 15, 19). Notably, actions are evaluated only in Line 9; at no other point are they explicitly evaluated. These modifications ensure that action evaluations are deferred as much as possible, providing sizeable computational savings.

Theorem 1. *If the estimator \hat{Q}_{init} is an underestimate, i.e., $\hat{Q}_{init}(b, a) \leq Q^*(b, a) \forall b \in \mathcal{B}, a \in \mathcal{A}$, where $Q^*(b, a)$ corresponds to the optimal Q -value, then both Lazy RTDP-Bel and Lazy LAO* will converge to the optimal policy.*

Proof Sketch. The idea is similar to the optimality proof for vanilla RTDP-Bel and LAO*. Since \hat{Q}_{init} is an underestimate, Bellman backups ensure $Q(b, a) \leq Q^*(b, a)$ for all belief-action pairs encountered throughout the search. Each backup progressively increases the value of the current best action (the action that minimizes Q -values) at belief states

along the current policy, bringing it closer to the optimal value. If a suboptimal action initially minimizes the Q -value at a belief state, updates eventually raise its value above that of the optimal action, ensuring the correct action is selected and the search converges to the optimal policy. \square

Full horizon laziness While the bottleneck in computing belief transitions is typically expensive transition or observation models, in some planning problems (as shown in Section 5), the bottleneck is verifying an action’s validity from a belief state. Here, successor beliefs can be efficiently computed, but determining action feasibility is computationally expensive. For example, in navigation tasks, an action may be deemed invalid from a belief b if executing it from any state s with $b(s) > 0$ could cause a collision. Verifying feasibility requires expensive collision checks, while computing successor beliefs assuming the action is valid is relatively cheap (as motion-primitives are known). In such cases, we can adopt what we call full-horizon lazy planning. Here, the policy is initially computed (using your favorite solver) assuming all actions are valid. Only actions in the computed policy undergo the expensive validation process. Invalid actions are marked infeasible, and the planner is queried again. This repeats until the policy consists entirely of valid actions. Similar to LazySP in deterministic settings, this approach defers the bottleneck operation until a complete policy is computed, minimizing the number of expensive queries. However, it is only applicable when successor beliefs can be computed efficiently and the primary computational cost lies in verifying action feasibility.

4.3 Q-value Estimators for Lazy Planning

Since the estimator \hat{Q}_{init} is only used to initialize the Q -values, it can be interpreted as an approximation of the initial Q -values, $Q_{init}(b, a)$, defined as

$$Q_{init}(b, a) = \mathcal{C}(b, a) + \sum_{z \in Z} P(z|b, a) \cdot \text{heur}(b_a^z). \quad (5)$$

A strong correlation between \hat{Q}_{init} and Q_{init} enhances the effectiveness of lazy planners. To guarantee optimality, \hat{Q}_{init} should underestimate Q_{init} (assuming the heuristic is admissible). However, a close approximation suffices for strong performance. Prior work in deterministic planning has explored using learned models to predict heuristic values (Veerapaneni, Saleem, and Likhachev 2023; Takahashi et al. 2019). We see potential in extending these ideas to stochastic settings. Domain-specific knowledge can also be leveraged to construct useful estimators.

Our focus is not on designing sophisticated estimators but on demonstrating the utility of lazy planners when paired with an estimator. In addition to the approaches mentioned, we present a simple domain-independent subsampling method that proved effective in our empirical analysis.

The key computational bottleneck in computing $Q_{init}(b, a)$ (Eq. 5) is determining successor beliefs b_a^z . This requires evaluating action a on belief b , which is expensive as b is a distribution over the state space \mathcal{S} . The action must be evaluated for every state s where $b(s) \neq 0$.³

³In many robotics problems, the state space is large. While the

Given the number of belief states typically encountered during search, the computational cost is significant.

To address this, we estimate $Q_{init}(b, a)$ using subsampling. We construct a discrete approximation \hat{b} of b by sampling states $s \sim b$ until $\text{supp}(\hat{b})$ reaches a predefined threshold Δ . The frequency of occurrence of each sampled state determines its probability in \hat{b} . Since $Q_{init}(\hat{b}, a)$ can be computed efficiently due to the reduced support and approximates $Q_{init}(b, a)$ well, we define the subsampled estimator as $\hat{Q}_{init}(b, a) = Q_{init}(\hat{b}, a)$.

The subsampling estimator in its vanilla form is effective when the heuristic satisfies $\text{heur}(b) = \mathbb{E}_{s \sim b} \text{heur}(s)$, a property common in robotics planning, especially in goal-directed tasks (e.g., an autonomous vehicle reaching a goal). However, if the heuristic does not satisfy this property—such as when it measures belief entropy (e.g., in active localization, where the cost of reducing uncertainty correlates well with the current uncertainty level (Spaan 2008))—correction terms can be introduced to mitigate bias. Although effective, the subsampled estimates are not guaranteed to be conservative. A detailed discussion of the subsampling estimators for our evaluated domains is included in the Appendix.

Additionally, in problem domains where belief transition costs stem primarily from the observation model, while querying the transition model is inexpensive, we propose an estimator inspired by Q^{MDP} heuristics (Hauskrecht 2000). Commonly used as admissible estimates in POMDP planning, Q^{MDP} heuristics assume the state becomes fully observable after a single action. We define this estimator as:

$$\hat{Q}_{init}^{MDP}(b, a) = \sum_{s \in \mathcal{S}} b(s) \left(c(s, a) + \sum_{s'} \mathcal{T}(s, a, s') \text{heur}(s') \right)$$

The estimator \hat{Q}_{init}^{MDP} is conservative if $\text{heur}(s')$ underestimates the cost of reaching the goal set G from s' in the underlying MDP defined by the transition model \mathcal{T} .

4.4 Discussion

The idea of lazily evaluating belief transitions by deferring their computation until necessary extends well beyond the two specific instantiations discussed above. It naturally applies to variants such as ILAO* (Hansen and Zilberstein 2001), IBLAO* (Warnquist, Kvarnström, and Doherty 2010), and AO* (Nilsson 2014). POMHDP (Kim, Salzman, and Likhachev 2019), which augments RTDP-Bel with multiple heuristics, can also incorporate laziness by using separate Q -estimators for each heuristic. Beyond these, the concept extends to other heuristic search methods like AEMS2 (Ross, Chaib-Draa et al. 2007), BI-POMDP (Washington 1997), and Satia and Lave (1973). With a few modifications, we also see potential in extending it to off-policy algorithms like AEMS (Ross, Chaib-Draa et al. 2007), which expand belief nodes outside the current best policy. Overall, lazy heuristic search for POMDP planning forms a broad class of algorithms. Finally, the core idea of deferring transition computations via Q -value estimation can be applied to MDPs as well.

support of a typical belief state during search is substantial, it still represents only a fraction of the total state space.




Figure 1: Planning times as a function of object pose uncertainty on the manipulation for pose estimation problem.

5 Results

The performance of lazy planners is evaluated across three problem domains: manipulation for pose estimation, indoor navigation, and outdoor rough terrain navigation. Without loss of generality, we represent the belief distribution as a discrete particle set and assume perfect observations.⁴ For all problem domains, the subsampling estimator $\hat{Q}_{init}(b, a) = Q_{init}(\hat{b}, a)$ samples \hat{b} to contain 15% of the total particles in b , i.e., Δ is set to 15% of the support of b . The reported planning time and cost statistics are averaged over runs where all algorithms solving at least 20% of problems have succeeded. Statistics for algorithms solving fewer than 20% are excluded to maintain representativeness.

5.1 Manipulation for Pose Estimation using Contacts

Planner	Success Rate	Planning Time	Cost
RTDP-Bel	0.87	307.49	0.69
Lazy RTDP-Bel	1.0	183.18	0.71
LAO*	0.93	253.67	0.69
Lazy LAO*	1.0	136.73	0.70

Table 1: Performance comparison on the manipulation for pose estimation problem.

In this domain, the goal is to localize a target object T (a charger plug) and complete an insertion task. The plug’s pose is not directly observable; instead, the robot manipulator infers it through binary contact observations (Fig. 2). This is useful when visual feedback is ineffective due to occlusions or poor lighting. The presence or absence of contact during actions reduces pose uncertainty. For instance, if no contact is observed while sweeping a volume, the object is absent in that region. Conversely, contact confirms the object’s presence in the swept volume.

Let the particle set $H = \{h_1, h_2, \dots\}$ represent the distribution of hypothesis poses for the target object T . The action space includes unit robot movements and compliant

⁴This is not a limitation; our framework can directly support other belief representations and noisy observation models.




Figure 2: The robot utilizing contact observations from different actions to reduce uncertainty about the pose of the object of interest, i.e., the charger port, before completing the insertion task.

controllers that track object surfaces upon contact. The observation space consists of the robot state q (from encoder readings) and a binary collision signal (from an F/T sensor). The belief state is given by $b = \{q, H\}$. The objective is to find a policy that fully localizes the target object ($|H| = 1$) while minimizing the robot’s expected travel distance. For a given object pose h_i and action a , computing the expected observation—i.e., the sequence of contact signals and robot states—requires forward simulating the action with the object model at h_i which involves complex mesh-to-mesh collision checks. Computing successor beliefs for $b = \{q, H\}$ requires simulating actions for all $h \in H$, making the process computationally expensive. Further details about the problem formulation are in Saleem, Veerapaneni, and Likhachev (2023, 2024). Since this is an active information-gathering problem, we define the belief state heuristic as a scaled version of the size of the set H (a measure of the belief entropy). We use the subsampling estimator with a minor modification (to ensure unbiased estimation), details about which can be found in the Appendix.

Results averaged over 100 runs are presented in Table 1. In each run, the planner resolves 3D positional uncertainties in the pose of T , with uncertainty along each axis randomly chosen between 4mm and 80mm. This uncertainty is discretized at a 2mm resolution to form the hypothesis set H , meaning a 30mm uncertainty per axis results in a start belief with 15^3 hypotheses. Each planner was given a 500-second timeout. The results show lazy planners outperform their vanilla counterparts by solving problems more efficiently through fewer belief transition evaluations. Lazy planners incurred roughly 50-60% of the planning times of their vanilla counterparts, enabling them to tackle harder problems and achieve higher success rates. Figure 1 shows planning times as a function of the number of hypothesis poses ($|H|$) in the start belief, i.e., start state uncertainty. The computational cost of computing transitions for a belief state is proportional to the number of hypotheses in it. As the start state uncertainty increases, the number of belief states with a high number of hypothesis poses encountered during the search also increases. As a result, the speedups offered by the lazy planners grow with the start state uncertainty.

5.2 Indoor Navigation

This is a 3D navigation domain, where the robot state is defined by the tuple (x, y, θ) . The map is known but the state of the robot is not directly observable, instead, the robot relies on a 1D LiDAR sensor mounted on it for perception. Computing expected observations from a given state require expensive raycasting operations to simulate the sensor. The belief over the robot’s state is represented as a set of parti-




Figure 3: Visualization of the indoor (left) and outdoor navigation environments (right).

cles $H = \{h_1, h_2, \dots\}$, each corresponding to a hypothesis robot pose h_i . Consequently, computing expected observations from a belief state requires performing raycasting from all hypothesis poses in the belief, making the belief transition computations expensive. The discrete action set \mathcal{A} consists of predefined motion primitives. We evaluate two variants of the problem. For both variants we use 3 indoor environments similar to the one we see in Fig 3.

Planner	Success Rate	Plan Time	Cost
RTDP-Bel	0.81	153.04	14.14
Lazy RTDP-Bel (\hat{Q}_{init})	0.90	32.46	14.26
Lazy RTDP-Bel (\hat{Q}_{init}^{MDP})	0.90	38.76	14.14
LAO*	0.71	185.37	14.14
Lazy LAO* (\hat{Q}_{init})	0.77	33.54	14.24
Lazy LAO* (\hat{Q}_{init}^{MDP})	0.77	40.43	14.14

Table 2: Performance comparison on the indoor navigation with stochastic transition dynamics problem.

With stochastic transition dynamics In this variant, uncertainty in the robot’s state arises from stochastic transitions. Certain regions in the map induce probabilistic motions, where the robot reaches the intended end state with probability 0.5, while slipping left or right with equal likelihood. The robot’s start state is known, and the objective is to reach a goal region G . The belief state is represented as a set of weighted particles, with the weight corresponding to the probability of each robot pose. Heuristic is the expected distance of the states in the belief from G , $heur(b) = \mathbb{E}_{s \sim b} \text{Dist}(s, G)$. Dist is the minimum distance of s from G on the MDP defined by the transition model \mathcal{T} , making the heuristic admissible. Alternatively, Dist can be the Euclidean distance between s and G , providing a quick

but inadmissible heuristic. We evaluate with the subsampled estimator, and since the computational expense of the belief transitions stems from the observation model, we also evaluate with the \hat{Q}_{init}^{MDP} estimator. Results presented in Table 2 are averaged over 200 runs, with random start states and goal regions. The timeout was 300 seconds. We observed that the lazy planners were particularly effective in this problem, improving planning speeds by 5x while computing policies with a marginally higher expected cost. The \hat{Q}_{init}^{MDP} estimator resulted in slightly higher planning times in comparison to the subsampling estimator, however, it is guaranteed to be a conservative estimate (while the subsampled estimator is not), resulting in solution costs identical to that of the vanilla planners.

Planner	Success Rate	Planning Time	Cost
Goal Directed Navigation			
RTDP-Bel	0.73	139.32	19.57
Lazy RTDP-Bel	0.89	84.91	19.61
LAO*	0.87	96.47	19.57
Lazy LAO*	0.99	53.42	19.60
Information Gathering			
RTDP-Bel	1.0	39.17	14.39
Lazy RTDP-Bel	1.0	21.32	14.25
LAO*	1.0	32.21	14.42
Lazy LAO*	1.0	16.75	14.87

Table 3: Performance comparison on the indoor navigation with start state uncertainty problem.

With start state uncertainty In this variant, uncertainty in the robot’s state arises from start state uncertainty, defined by an initial set of equally likely hypothesis poses. The transition dynamics is deterministic, allowing the belief to be represented as a set of unweighted particles. This problem has two versions: (1) the robot must localize itself, and (2) the robot must reach a goal region G . For the information-gathering version, the goal belief, belief heuristic, and Q -estimator are defined identically to the manipulation problem. In the goal-directed version, they follow the definitions used for stochastic transition dynamics. Results for both versions, averaged over 200 runs, are presented in Table 3, with each run using a random set of 100 to 250 start hypothesis poses. Similar to the manipulation domain, the lazy planners on average took 60% of the planning time of their vanilla counterparts at the expense of marginally higher costs.

5.3 Rough Terrain Navigation

This is a 3D outdoor navigation problem with deterministic transitions and observations. Uncertainty in the robot’s state arises from start state uncertainty, with the belief represented as a set of unweighted particles. Unlike indoor navigation, the robot lacks a 1D LiDAR sensor; instead, the environment contains landmarks that the robot can observe if within a predefined distance, making the observation model efficient to query. Consequently, both the transition and observation models are computationally cheap. However, in this domain,

Planner	Success Rate	Planning Time	Cost
Goal Directed Navigation			
RTDP-Bel	0.01	-	-
Lazy RTDP-Bel	0.21	1297.33	12.66
FH Lazy RTDP-Bel	0.90	109.9	12.66
LAO*	0.09	-	-
Lazy LAO*	0.33	886.64	12.66
FH Lazy LAO*	0.90	139.52	12.66
Information Gathering			
RTDP-Bel	0.07	-	-
Lazy RTDP-Bel	0.19	1015.51	14.58
FH Lazy RTDP-Bel	0.83	110.38	14.58
LAO*	0.09	-	-
Lazy LAO*	0.26	642.78	14.58
FH Lazy LAO*	0.81	163.79	14.58

Table 4: Comparison on the outdoor navigation problem.

an action is feasible from a belief state only if it can be accurately tracked from all hypothesis states in the belief. Due to rough terrain, certain states may cause the robot to get stuck in pits, fail to climb inclines, or execute motions inaccurately. Hence, validating an action requires expensive roll-outs in a physics simulator, making action verification the primary bottleneck. Since the expected trajectory that will be tracked from the hypothesis states if the action were valid is known, and the observation model is efficient, belief transitions can be computed quickly. The bottleneck operation is simply the verification of the validity of actions.

Similar to the indoor navigation case, we evaluate two versions of the problem. Table 4 presents the performance averaged over 100 runs, with each run containing 30–50 randomly sampled start state hypothesis and a timeout of 2400 seconds. FH Lazy RTDP-Bel and FH Lazy LAO* correspond to the full-horizon lazy variants discussed in Section 4.2. We observe that the full-horizon lazy planners demonstrate strong performance boosting success rates by close to 90% in comparison to the vanilla counterparts. By deferring action evaluations until after policy computation, the full-horizon lazy planners achieve a 10x speed up over the 1-step lazy planners. The high computational expense of the physics simulations make the impact of laziness pronounced in this domain, highlighting the utility of lazy planners.

6 Conclusion

In this work, we introduced Lazy RTDP-Bel and Lazy LAO*, two specific instances of a broader class of lazy heuristic search solvers for POMDPs with expensive-to-compute belief state transitions. These algorithms leverage Q -value estimates to defer expensive belief transition evaluations until they are truly necessary. By postponing these computations, our approach significantly enhances planning efficiency while preserving solution quality. We validated our methods across three challenging robotics problems, achieving substantial speedups over conventional solvers in all of them.

Acknowledgements

This work was supported in part by Honda Research Institute and in part by the ARO-sponsored DURIP grant W911NF-21-1-0050.

References

- Barto, A. G.; Bradtke, S. J.; and Singh, S. P. 1995. Learning to act using real-time dynamic programming. *Artificial intelligence*, 72(1-2): 81–138.
- Bertsekas, D. 2012. *Dynamic programming and optimal control: Volume I*, volume 4. Athena scientific.
- Bhardwaj, M.; Choudhury, S.; Boots, B.; and Srinivasa, S. 2021. Leveraging experience in lazy search. *Autonomous Robots*, 45(7): 979–996.
- Bonet, B.; and Geffner, H. 2009. Solving POMDPs: RTDP-Bel vs. Point-based Algorithms. In *IJCAI*, 1641–1646.
- Brethel, S.; Gindele, T.; and Dillmann, R. 2014. Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs. In *17th international IEEE conference on intelligent transportation systems (ITSC)*.
- Cohen, B.; Phillips, M.; and Likhachev, M. 2015. Planning single-arm manipulations with n-arm robots. In *Proceedings of the International Symposium on Combinatorial Search*, volume 6, 226–227.
- Dellin, C.; and Srinivasa, S. 2016. A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors. In *Proceedings of the international conference on automated planning and scheduling*, volume 26, 459–467.
- Hansen, E. A.; and Zilberstein, S. 2001. LAO*: A heuristic search algorithm that finds solutions with loops. *Artificial Intelligence*, 129(1-2): 35–62.
- Hauskrecht, M. 2000. Value-function approximations for partially observable Markov decision processes. *Journal of artificial intelligence research*, 13: 33–94.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.
- Kim, S.-K.; Salzman, O.; and Likhachev, M. 2019. POMHDP: Search-based belief space planning using multiple heuristics. In *Proceedings of the international conference on automated planning and scheduling*, volume 29.
- Kurniawati, H.; Hsu, D.; and Lee, W. S. 2008. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, volume 2008. Citeseer.
- Kurniawati, H.; and Yadav, V. 2016. An online POMDP solver for uncertainty planning in dynamic environment. In *Robotics Research: The 16th International Symposium ISRR*, 611–629. Springer.
- Lee, J.; Kim, G.-H.; Poupart, P.; and Kim, K.-E. 2018. Monte-Carlo tree search for constrained POMDPs. *Advances in Neural Information Processing Systems*, 31.
- Mandalika, A.; Choudhury, S.; Salzman, O.; and Srinivasa, S. 2019. Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, 745–753.
- Mandalika, A.; Salzman, O.; and Srinivasa, S. 2018. Lazy receding horizon A* for efficient path planning in graphs with expensive-to-evaluate edges. In *Proceedings of the international conference on automated planning and scheduling*, volume 28, 476–484.
- Nilsson, N. J. 2014. *Principles of artificial intelligence*. Morgan Kaufmann.
- Pajarinen, J.; and Kyrki, V. 2017. Robotic manipulation of multiple objects as a POMDP. *Artificial Intelligence*, 247.
- Pineau, J.; Gordon, G.; Thrun, S.; et al. 2003. Point-based value iteration: An anytime algorithm for POMDPs. In *Ijcai*, volume 3, 1025–1032.
- Ross, S.; Chaib-Draa, B.; et al. 2007. AEMS: An anytime online search algorithm for approximate policy refinement in large POMDPs. In *IJCAI*, 2592–2598.
- Saleem, M. S.; Veerapaneni, R.; and Likhachev, M. 2023. Preprocessing-based planning for utilizing contacts in semi-structured high-precision insertion tasks. *IEEE Robotics and Automation Letters*, 8(11): 6947–6954.
- Saleem, M. S.; Veerapaneni, R.; and Likhachev, M. 2024. A POMDP-based hierarchical planning framework for manipulation under pose uncertainty. *arXiv preprint arXiv:2409.18775*.
- Satia, J.; and Lave, R. 1973. Markovian decision processes with probabilistic observation of states. *Management Science*, 20(1): 1–13.
- Shani, G.; Pineau, J.; and Kaplow, R. 2013. A survey of point-based POMDP solvers. *Autonomous Agents and Multi-Agent Systems*, 27: 1–51.
- Silver, D.; and Veness, J. 2010. Monte-Carlo planning in large POMDPs. *Advances in neural information processing systems*, 23.
- Smith, T.; and Simmons, R. 2012. Heuristic search value iteration for POMDPs. *arXiv preprint arXiv:1207.4166*.
- Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. DESPOT: Online POMDP planning with regularization. *Advances in neural information processing systems*, 26.
- Spaan, M. T. 2008. Cooperative active perception using POMDPs. In *AAAI 2008 workshop on advancements in POMDP solvers*.
- Takahashi, T.; Sun, H.; Tian, D.; and Wang, Y. 2019. Learning heuristic functions for mobile robot path planning using deep neural networks. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, 764–772.
- Veerapaneni, R.; Saleem, M. S.; and Likhachev, M. 2023. Learning local heuristics for search-based navigation planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, 634–638.
- Warnquist, H.; Kvarnström, J.; and Doherty, P. 2010. Iterative Bounding LAO. In *ECAI 2010*, 341–346. IOS Press.
- Washington, R. 1997. BI-POMDP: Bounded, incremental partially-observable Markov-model planning. In *European Conference on Planning*, 440–451. Springer.