

Heuristics for Partially Observable Stochastic Contingent Planning

Guy Shani*

Software and Information Systems Engineering, Ben Gurion University, eBay Research
ORCID (Guy Shani): <https://orcid.org/0000-0002-3655-7148>

Abstract. Acting to complete tasks in stochastic partially observable domains is an important problem in artificial intelligence, and is often formulated as a goal-based POMDP. Goal-based POMDPs can be solved using the RTDP-BEL algorithm, that operates by running forward trajectories from the initial belief to the goal. These trajectories can be guided by a heuristic, and more accurate heuristics can result in significantly faster convergence. In this paper, we develop a heuristic function that leverages the structured representation of domain models. We compute, in a relaxed space, a plan to achieve the goal, while taking into account the value of information, as well as the stochastic effects. We provide experiments showing that while our heuristic is slower to compute, it requires an order of magnitude less trajectories before convergence. Overall, it thus speeds up RTDP-BEL, particularly in problems where significant information gathering is needed.

1 Introduction

Autonomous agents acting to achieve their goals in stochastic environments often face partial observability, where important information for completing the task is hidden, but can be sensed. The agent must consider the value of information, that is, whether it is better, in terms of expected cost, to invest effort in sensing for the hidden information, or to try and reach the goal without the information.

Consider, for example, a robot navigating in an uncertain environment, where the ground may be unstable. The robot can either operate a sensor that examines the ground to check its stability, later allowing it to move rapidly only over stable areas, or move slowly and carefully without operating the sensor first. The choice as to whether it is cost effective to use the sensor depends on the probability of unstable areas, the time it takes for the sensor to scan the ground, and the cost of moving through possibly unstable areas. This tradeoff is known as the value of information.

Partially observable Markov decision processes (POMDPs) [32] are a mathematical model designed for properly estimating the value of information over long planning horizons in such environments. Goal-POMDPs [16] are a subclass where the agent must achieve a goal, upon which the execution terminates. POMDPs can be either specified using a "flat" representation, where each state is a unique entity, or using a factored representation, where a state is composed of state variables. We focus here on a particular factored representation using an extension of the PDDL (Planning Domain Description Language) representation which is popular in the automated planning

literature [17]. PDDL defines actions by preconditions and effects, allowing for a compact representation of complex problems.

POMDP solvers often use the concept of a belief state, a distribution over the currently possible world states, and compute a value function, mapping a belief state b to the expected cost to the goal from that belief. The RTDP-BEL (Real-Time Dynamic Programming in Belief Space) algorithm [11] computes a value function by running trajectories in belief space from the initial belief to a goal. It depends on a heuristic to guide the trajectories towards goal states. Good trajectories that visit belief states that are typically visited on an optimal path to the goal can lead to much faster convergence. We develop here heuristics that steer RTDP-BEL towards such good trajectories.

An advantage of using a PDDL representation is that it paves the way to employing the multitude of heuristics developed for various automated planning solvers. For example, the delete relaxation approach [6], that ignores negative effects, can be used to rapidly estimate the distance to the goal using techniques such as h_{max} [10] and h_{ff} [19]. This paper suggests how these techniques can be used for generating heuristic cost-to-go estimations for belief states.

More specifically, during the heuristic computation, we show how in the forward exploration in relaxed space we can take sensing actions into account, considering only actions whose preconditions hold in all possible states. Thus, we can compute an heuristic estimation of the information needed for achieving a particular plan. This results in more accurate cost-to-go heuristic estimations.

We suggest here heuristics and test them in the context of the RDTP-BEL offline POMDP solver, which has shown superior performance compared to the popular point-based approaches in goal-based domains [11]. We hence do not evaluate our methods here in the context of point-based solvers. Although online POMDP solvers [29] have gained much popularity, offline planning has a number of advantages, mainly low online computational burden, allowing the previously computed policy to be rapidly employed in weak end devices. Our empirical evaluation compares our heuristics to popular methods, such as the well known Q_{MDP} heuristic [9, 18], and focusing only on the most likely state. We show that our heuristics, while slower to compute, improves the convergence of RTDP-BEL such that the overall runtime is significantly reduced, especially in tasks that require costly information gathering action sequences.

2 Background

A goal Markov decision process (goal-MDP) [3, 5] is a tuple $\langle S, S_0, A, tr, G, C \rangle$ where S is a state space, A is an action, tr :

* Corresponding Author. Email: shanigu@bgu.ac.il, gshani@ebay.com.

$S \times A \times S \rightarrow [0, 1]$ is a transition function. $tr(s, a, s')$ is the probability of executing action a at state s and reaching state s' . $S_0 \subset S$ is a set of possible initial states, and $G \subset S$ is a set of goal states, that the agent must reach. $C : S \times A \rightarrow \mathbb{R}$ is a cost function. $C(s, a)$ is the cost for applying action a at state s . In goal-MDPs for $s \in G$, $C(s, a) = 0$ and $C(s, a) > 0$ otherwise.

A solution to a goal MDP is a policy $\pi : S \rightarrow A$ that assigns an action to every state. One can create a policy by computing a value function $V : S \rightarrow \mathbb{R}$ assigning a value for each state. This value function is optimal if $V(s)$ is the minimal expected cost of reaching a goal state from s . The Bellman backup:

$$V(s) \leftarrow \min_{a \in A} C(s, a) + \sum_{s' \in S} tr(s, a, s') V(s') \quad (1)$$

can be used to iteratively update the value function. Value iteration applies the Bellman update over all states repeatedly, and converges to the optimal value function.

Real Time Dynamic Programming (RTDP) [2] applies the Bellman update only along trajectories that start at a state in S_0 and terminate at a state in G . Thus, RTDP is able to focus on states that are visited along an optimal path to the goal, avoiding repeated Bellman backups over other states, that are less important. This allows RTDP to converge much faster than value iteration in many domains.

A goal-POMDP is a tuple $\langle S, S_0, b_0, A, tr, G, C, \Omega, O \rangle$ where $\langle S, S_0, A, tr, G, C \rangle$ is a goal-MDP, known as the underlying MDP. b_0 is a distribution over the possible initial states in S_0 , Ω is a set of possible observations, and $O(a, s, o)$ is the probability of observing o after applying action a and reaching state s . In this paper we focus on deterministic observations, and hence, $O(a, s, o) \in [0, 1]$. This is not truly a limitation, because every POMDP with stochastic observations can be translated into a deterministic observation POMDP, by adding a state variable whose value changes stochastically, yet observed deterministically. Consider for example a noisy wall detection sensor. Instead of observing a wall with stochastic errors, we can add a red light to the sensor, which serves as the state variable. The red light is stochastically lit by the sensor when there the sensor detects a wall, but the agent observes whether the light is on deterministically.

A belief state is a distribution over states, where $b(s)$ is the probability of state s . Given a belief b an action a and an observation o one can compute the new belief state b_a^o using:

$$b_a^o(s') = \begin{cases} \kappa \cdot \sum_s b(s) tr(s, a, s') : O(a, s', o) = 1 \\ 0 : O(a, s', o) = 0 \end{cases} \quad (2)$$

where κ is a normalizing factor. A popular technique for solving POMDPs is by computing a value function over the belief space MDP, defined over the all belief states. While this MDP has an infinite amount of states, its value function is convex, which allows for a convenient representation using a set of half-spaces called α -vectors.

In this paper, however, we do not use α -vectors. Instead, we use RTDP in belief space, called RTDP-BEL [16, 9, 11] to focus the exploration of the belief space to beliefs that are likely to occur along an optimal path to the goal. RTDP-BEL is presented in algorithm 1. Lines 6-12 produce a forward trajectory. The values along that trajectory are updated when moving forward, and thus, one can show that the trajectory must terminate at a goal state. It is useful, however, to also update the states after the trajectory has terminated moving backward from the goal (lines 13-5).

In the value function update (lines 9,15) when $V(b)$ was not yet computed, we can use a cost-to-go heuristic, to initialize it. If this heuristic is well correlated with the optimal value of $V(b)$, it can

Algorithm 1: RTDP-BEL

```

1  RTDP-BEL
2  while  $V$  has not converged do
3       $b \leftarrow b_0$ 
4       $s \leftarrow$  sample from  $b_0$ 
5       $l \leftarrow$  the empty list
6      while  $s \notin G$  do
7          add  $b$  to  $l$ 
8           $a^* \leftarrow \operatorname{argmin}_a C(b, a) + \sum_o pr(o|a, b) V(b_a^o)$ 
9           $V(b) \leftarrow C(b, a^*) + \sum_o pr(o|b, a^*) V(b_a^o)$ 
10          $s' \leftarrow$  sample from  $tr(s, a^*, \cdot)$ 
11          $o \leftarrow$  the observation s.t.  $O(a, s', o) = 1$ 
12          $b \leftarrow b_a^o$ 
13     foreach  $b \in l$  in reversed order do
14         foreach  $a \in A$  do
15              $a^* \leftarrow \operatorname{argmin}_a \sum_o pr(o|a, b) V(b_a^o)$ 
16              $V(b) \leftarrow C(b, a^*) + \sum_o pr(o|b, a^*) V(b_a^o)$ 

```

steer RTDP-BEL to chose good actions in line 8 that lead it faster to the goal. The heuristic must be optimal (a lower bound) over the optimal costs, otherwise, when an action may seem optimal, RTDP-BEL will stop exploring alternative actions from that belief, and may converge to a sub-optimal value function. As we shall later see, this is a real problem for some heuristics. Our paper focuses on developing strong heuristics for initializing $V(b)$.

Given a value function $V : B \rightarrow \mathbb{R}$ computed by RTDP-BEL, we can define a policy using:

$$\pi_V(b) = \operatorname{argmin}_a C(b, a) + \sum_o pr(o|b, a) V(b_a^o) \quad (3)$$

Unlike value functions represented using α -vectors, a value function represented using a direct mapping from belief states to values, does not generalize to other belief states. Thus, when computing the policy for a belief b , it might be that some values $V(b_a^o)$ are missing. In this case, we can use the heuristic again for these missing values.

POMDPs can also be described in a structured manner, following conventions developed in the planning community, specifically for stochastic planning and contingent planning with sensing actions [1]. We now provide such a specification. A stochastic, partially observable, contingent planning problem is a tuple (overloading previous definitions) $\langle P, A, G, I \rangle$. P is a set of facts. Each state is a truth assignment for all facts.

A is a set of actions. Each action $a \in A$ is a tuple $\langle pre, eff, obs \rangle$ where pre is the precondition, a set of facts that must hold before applying the action. eff is a set of conditional effects of the form (c, e) where c is a set of facts that form the condition — if all facts in c hold then the condition is fulfilled. e is a probabilistic formula — a set (e_1, \dots, e_k) of possible probabilistic effects, where e_i is a set of facts. For each fact l in e_i , if l was false prior to the execution of the action, it would become true with probability $pr(e_i)$ after the action is executed. obs is a set of facts whose value is observed in the new state following the action. Action have a cost $C(a)$, $C(s, a) = C(a)$, $s \notin G$, and $C(s, a) = 0$, $s \in G$.

We assume here, as is often done in automated planning literature, that all preconditions and conditions contain only positive facts. This is not truly a limitation, as one can always maintain for every fact p an auxiliary fact np , ensuring that $p = \neg np$ always holds.

I is a set of probabilistic formulas over the facts, defining the possible initial values of facts, thus also defining the initial belief state. G is a set of goal facts.

The agent operates in belief space. Hence, before applying an action a in a belief b , the agent must ensure that the preconditions $a.pre$ hold in every state s such that $b(s) > 0$. This approach is useful for

limiting the branching factor — the amount of actions that must be considered at every belief state. As we assume here deterministic observations, each observation effectively eliminates from the next belief every state that does not agree with the observed value of one of the facts in $a.obs$.

3 Related Work

In this paper we use our heuristics within the RTDP-BEL algorithm. A popular alternative to RTDP-BEL is the point-based approach, where a value function in the form of a set of α -vectors is computed over a finite set of belief states [28, 30, 24, 23, 33, 25]. While RTDP-BEL is well suited for goal-POMDPs, point-based methods may also work well in such domains. It is quite possible that our heuristics can be helpful in obtaining good belief states for point-based approaches as well. That is, one may be able to design a new point-based algorithm that uses our methods as a heuristic for collecting belief points. We leave this for future research.

Another approach for acting in partially observable environments is through online approaches, such as POMCP [29] and DESPOT [31]. Online algorithms expand a belief tree from the current belief state, and then choose an action. When constructing the tree, heuristics are very useful. However, in many of the domains that we experiment with, the number of possible actions at each belief is huge. This renders online algorithms practically useless, because with a huge branching factor, the trees grow beyond manageable sizes. Attempts at using such heuristics in POMCP over similar domains did not scale beyond tiny problems [7, 8].

Bertoli et al. [4] consider similar motivations to our approach when developing heuristics for search in belief space. They capture knowledge at a higher degree than the state based level, estimating the probability that some formulas hold. They estimate a value for sensing actions only at a myopic level, which RTDP-BEL already handles.

Bryce et al. [14, 15] extend methods for computing heuristics in classical planning to belief space. They use the planning graph structure, which is a more elaborate construct than the simple layered graph used for computing h_{ff} . To adapt the planning graph to belief space, they show how to compute a heuristic estimate for multiple states together. In this context, however, sensing action provide no value. We evaluate this idea, although not the efficient computation, in our experiments below, showing that it does not handle well domains requiring information gathering.

Wang and Dearden [35] also provide a myopic value of information estimation, estimating the long range effects of an immediate sensing action. They do not, however, take into account sensing actions that will be needed later in the plan.

Our belief maintenance method is simple, maintaining the distribution over all states that have a non-zero probability for each belief. In many domains, the number of non-zero probability states is not very large, so this representation is not very costly. We are also caching state transitions once they were computed to avoid repeated computations. Our attempts at using more sophisticated methods (e.g. [13]) did not show any advantage. In any case, the bottleneck in our algorithms is the heuristic computation, not the belief maintenance. That being said, many methods were suggested for more efficient representations [12, 34], and we leave the examination of such methods in the context of RTDP-BEL for future research.

Kim et al. [21] also investigate the use of heuristics in RTDP-BEL. They suggest to use multiple heuristics together, one admissible, and some non admissible, but more accurate. They show that they can still guarantee convergence, while taking into account the more accu-

rate heuristic. In their experiments, they use domain-specific heuristics, while we focus here on domain independent heuristics. One may possibly aggregate some of the heuristics that we evaluate, e.g. our belief-based, state-based, and Q_{MDP} together using their methods.

Saborio and Hertzberg [27] use subgoals to construct an heuristic estimate in robotics environments. They compute problem features, and evaluate how many features are satisfied in a given state. They do not, hence, estimate the value of information in sensing actions that require lengthy action sequences to execute from the current belief. Subgoals were investigated in the classical planning literature, e.g., through the computation of landmarks [26], but to date have not been able to substantially outperform other heuristics. We leave the investigation of subgoals in the context of RTDP-BEL for future research.

4 Heuristics In Belief Space

We now describe our main contribution — a heuristic in belief space that takes into account the need for sensing actions. Most, if not all, previous heuristics focus on computing a heuristic for each state, and then aggregating these heuristic values. We, however, suggest a method that directly computes a plan in a relaxed belief space.

Our methods currently assume unit costs. Extensions of our heuristics to non-unit costs is left for future research.

4.1 State-based Heuristics

Initializing a heuristic optimistic bound for POMDPs was previously suggested. Perhaps the most popular methods are based on the value function for the underlying MDP [18]. Given a computed value function V_{MDP} one can define the heuristic value for a belief state by either focusing on the most likely state: $s_{ML} = \operatorname{argmax}_s b(s)$, $h_{ML}(b) = V(s_{ML})$, or the Q_{MDP} heuristic that uses the weighted sum: $h_{QMDP}(b) = \sum_s b(s)V(s)$. More elaborate attempts, such as FIB [18], did not show a substantial improvement over Q_{MDP} .

In the planning community, there are many possible heuristics for estimating the cost of reaching the goal. A popular choice is to compute distances in a relaxed problem, known as delete-relaxation [19]. In delete-relaxation we ignore negative facts in the effects of actions. Thus, once a fact is added to a state, no action can remove it.

We can hence compute, for a given state, layers of positive facts iteratively (Algorithm 2). The first layer F_0 contains all the positive facts that hold in the input state s (line 2). Then, we apply all actions that can be executed at F_0 (line 7), and their positive effects are added to the next layer F_1 (line 8). This is continued until no new facts can be added, or until all goal facts appear in F_i (line 9). Then, we can compute a heuristic estimate. For example, we can return the number of layers until all goal facts were achieved (lines 12-13), known as the h_{max} heuristic, or sum the indexes of the layers where each goal fact appeared for the first time, known as the h_{add} heuristic [10]. However, h_{add} is inadmissible, as it does not acknowledge that several facts can be achieved together. For RTDP, this can be a serious disadvantage, as inadmissible heuristics cause RTDP to converge to a sub optimal solution.

The implementation shown in Algorithm 2 is simplistic, and there are several ways to improve it. For example, one can add to F_i only new facts that did not appear in some F_j , $j < i$. Then, we can check for the next level only actions that have some precondition in F_j , because there is no reason to execute an action twice. This is inaccurate for conditional effects that may have been inactive in past layers, and are now becoming active. However, we can take each action with conditional effects and create a set of actions, one for each condition.

Algorithm 2: Classical Delete Relaxation Heuristic

```

1  $h_{max}(s)$ 
2  $F_0 \leftarrow$  all positive facts in  $s$ 
3  $i \leftarrow 0$ 
4  $done \leftarrow false$ 
5 while  $\neg done$  do
6    $F_{i+1} \leftarrow F_i$ 
7   foreach  $a \in A$  s.t.  $a.pre$  hold in  $F_i$  do
8     Apply  $a$  in  $F_i$ , add all positive effects to  $F_{i+1}$ 
9     if  $G$  holds in  $F_i$  or  $F_{i+1} = F_i$  then
10        $done \leftarrow true$ 
11    $i \leftarrow i + 1$ 
12 if  $G$  holds in  $F_i$  then
13   return  $i$ 
14 return  $\infty$ 

```

As in delete relaxation actions are executed in parallel, this is equivalent to executing the original action with multiple conditions activated. Also, as we repeatedly estimate a heuristic for a state s from different belief states, we can cache the state estimations. In many domains, all possible states become cached, and hence, although new beliefs are constantly generated, no new heuristic computations are needed. We use these, and other, less important, modifications to expedite the computation, but leave them out of the pseudo-code for ease of exposition.

Bryce et al. [15] suggested to extend this method to belief states, by computing, e.g., h_{max} for each possible state. Then, a possible heuristic can be the weighted sum of heuristic values: $h_{max}(b) = \sum_s b(s)h_{max}(s)$. They show how this computation can be done efficiently without handling each state independently, but in our implementation we used a simple approach, computing h_{max} for each state and caching results for future computations.

An important disadvantage of h_{max} is that it ignores actions that are executed in parallel at the same layer. Thus, it often grossly underestimate the cost to achieving the goal. The h_{ff} heuristic further improves upon h_{max} in computing a plan in delete-relaxation space. This plan can be computed by maintaining, for each fact l , the action that first generated it. Upon reaching the goal, we trace back, identifying for each goal fact the action that generated it, and then continuing backwards to the precondition facts of that action, until all preconditions are satisfied at the input state. Then, the length of the plan becomes the heuristic estimate for that state. h_{max} is admissible in the stochastic sense, i.e. the value it computes is an optimistic estimate of the real expected cost. h_{ff} however is inadmissible, but in the domains that we experiment with, it does not over estimate.

In all the above heuristics, actions that only observe facts, but have no effect, are completely ignored. That is, the heuristic value of observations, and hence, the value of information, cannot be computed. This does not mean, though, that RTDP-BEL ignores these actions. As it iterates over all actions, and over all observations, it computes the myopic value of information of immediate actions. However, these heuristics provide to value to sensing actions that are distant from the current position of the agent.

4.2 Supporting Stochastic Effects

Delete-relaxation heuristics were originally designed for deterministic domains, and thus ignore stochastic effects. We suggest the following straight forward method for handling stochastic effects. For a fact l identified at level i that is a stochastic effect of an action with probability $pr(l)$, we insert the fact only at level F_{i+c} , where $c = \lceil \frac{1}{pr(l)} \rceil$. That is, we assume that the fact would only be achieved

Algorithm 3: Belief Delete Relaxation Heuristic

```

1  $h_{max}(b)$ 
2 Let  $s_{ML}$  be a most likely state in  $b$ 
3  $Valid = \{s : b(s) > 0\}$ 
4 foreach  $s \in Valid$  do
5    $F_0^s \leftarrow$  all positive facts in  $s$ 
6    $i \leftarrow 0$ 
7    $done \leftarrow false$ 
8   while  $\neg done$  do
9     foreach  $s \in Valid$  do
10        $F_{i+1}^s \leftarrow F_i^s$ 
11       foreach  $a \in A$  s.t.  $a.pre$  hold in  $\cap_{s \in Valid} F_i^s$  do
12         foreach  $s \in Valid$  do
13           Apply  $a$  in  $F_i^s$ , add all positive effects to  $F_{i+1}^s$ 
14           if  $a$  is a sensing action then
15             if  $\exists p \in a.obs : (p \in F_i^s, p \notin F_i^{s_{ML}}) \vee (p \notin F_i^s, p \in F_i^{s_{ML}})$  then
16               Remove  $s$  from  $Valid$ 
17           if  $G$  holds in  $\cap_{s \in Valid} F_i^s$  or  $\forall s \in Valid, F_{i+1}^s = F_i^s$  then
18              $done \leftarrow true$ 
19            $i \leftarrow i + 1$ 
20 if  $G$  holds in  $\cap_{s \in Valid} F_i^s$  then
21   return  $i$ 
22 return  $\infty$ 

```

after c executions of the action. This is equivalent to the expected number of times that a must be executed before l is achieved.

When computing our delete relaxation heuristics, we wait until the level where the fact is added, before deciding which action added it. That is, there may be another action that achieved this fact, deterministically or with higher probability, earlier.

4.3 Belief-based Heuristic

We now describe a new heuristic that is based on the delete-relaxation approach above, but extends it into belief space, allowing us to assign a value for sensing actions. The main difference from the classical method is in maintaining a set of achieved facts for each possible state. An action can be executed only if all states agree that it's preconditions hold. The heuristic computation must reach the goal given all possible states, or identify invalid states — states that were observed to be inconsistent with the true world state.

We choose one state, e.g. a maximum likelihood state, and compute the heuristic as if it is the real current state. Then, we can eliminate from considerations other states that do not agree with the assumed real state about the value of some observation. As the observations in our model are deterministic, each observation may eliminate some states, and thus may help us in enabling additional actions whose precondition was not met in the eliminated states. This approach allows us to evaluate which sensing actions provide useful information that allows us to identify the true world state.

Algorithm 3 presents the pseudo-code for this heuristic. We begin by choosing a most likely state, denoted s_{ML} to serve as the assumed real world state (line 2). We then define a set $Valid$ initialized to all states with non-zero probability in the current belief state b . We maintain for each valid state s the set of positive facts F_i^s (lines 4-5). We then begin computing the layers of positive facts for all states.

We iterate over all actions, whose precondition holds for all currently valid states. That is, we consider only facts that appear in the intersection of all valid F_i^s (line 13). Then, for all valid states, we update F_{i+1}^s . Due to conditional actions, different states may acquire different new facts. For ease of exposition, we are not showing in the pseudo-code the treatment of stochastic effects (Section 4.2).

We now consider sensing actions (lines 14-16). For each such action a that observes a set $a.obs$ of facts, we check if there is a fact p with different values in $F_i^{s_{ML}}$ and F_i^s (line 15). If this is the case, then, assuming s_{ML} is the true state of the system, s must be impossible, as observations are deterministic. Thus s is removed from the set *Valid* of currently valid states (line 16).

This process ensures that, just as the sets F_i^s only grow with additional layers, the set *Valid* can only become smaller. In many cases *Valid* eventually consists only of s_{ML} and the rest of the layers are computed in a similar manner to Algorithm 2. We also use the approach suggested above for handling stochastic effects.

For computing a plan in relaxed space for the h_{ff} heuristic, we must again keep track of which action first generated a fact l . We consider only facts that hold for all valid states. Such a new fact l can be introduced in a number of ways; First, an action a can add an effect l to all states, due to unconditional effects. In that case we mark a as the action that generated l .

Second, an action can add l to some valid states due to conditional effects. Then, it may be that several actions add l for different valid states until for all valid states F_i^s contains l . In that case we mark only the last action that added l as the action that generated l . This is imprecise, because it might be that many actions are needed for different valid states. However, this is an optimistic assumption, and thus it cannot cause the heuristic to become non admissible.

Third, it may be that due to an observation, several states have become invalid, and l holds in all remaining valid states. In that case we attribute the achievement of l to the sensing action. Again, this is an optimistic assumption, because it might be that a set of sensing actions, each eliminating some states, was required to ensure that l holds in all current states. As above, such optimistic assumptions cannot cause the heuristic to become non admissible.

We now reconstruct the delete-relaxation plan backwards. Then, we return the number of actions in the plan as the heuristic estimate.

Caching of heuristic computations here is more difficult, but still useful. As $h(b)$ does not depend on the distribution of b , we cache the results given s_{ML} and the set of states $\{s : b(s) > 0\}$. Hence, if we encounter another belief state that has the same most likely state and the same set of possible states, then we can use the cached heuristic computation. This is much less effective than caching individual state heuristic estimations. However, in stochastic domains, we often encounter beliefs with the same possible states, but with slightly different probabilities. In such cases, caching is very effective.

Our belief delete relaxation heuristics use the same stochastic effects mechanism described earlier as the state-based approaches.

Complexity: Computing the delete relaxation graph requires $O(|A|^2)$ operations, assuming a bounded number of action effects, as any action can be executed only once. Hence, computing our heuristics for a belief state b requires $O(|b||A|^2)$ operations, where $|b|$ is the number of states where $b(s) > 0$. This is needed for every heuristic estimate computation. Classical heuristics such as Q_{MDP} require solving the underlying MDP, which is in theory polynomial (using linear programming), but given the more practical value iteration requires $O(|S|^2|A|k)$ operations, where k is the number of iterations. Given the value function, computing a heuristic estimation requires $O(|b||A|)$ operations.

5 Empirical Evaluation

We now report an empirical evaluation, comparing the performance of the various heuristics that we suggest, in the context of RTDP-BEL. All methods were implemented in C# within the unified

planning framework¹. The experiments were conducted on an i7 3.40GHz CPU with 8 cores, and 64GB RAM.

Methods: We compare the following heuristics; First, we use two MDP-based heuristics — Q_{MDP} and the value of the most likely state (Section 4.1). These heuristics require the computation of a value function for the underlying MDP. In some of our problems, simple value iteration did not converge in under 5 minutes. Using regular RTDP over the state space, on the other hand, is insufficient. This is because RTDP focuses on states that are visited on good paths to the goal in the MDP. However, in the POMDP, there are important states that do not belong to any such path. This is most pronounced in cases where information gathering is needed. We hence use the following compromise — we compute initially an MDP value function using RTDP, and later, during the execution of RTDP-BEL, when reaching a state that was not visited by the initial RTDP execution, we run RTDP again for 100 additional iterations starting at that state. When reporting runtime, we exclude the initial RTDP execution, as it is not optimized, and other value function methods may be faster. Below we denote the two MDP based methods as *ML* for most likely state, and Q_{MDP} .

We use two heuristics based on h_{ff} that we describe above — computing the expected cost over the possible states using $h_{ff}(s)$ (Section 4.1), and computing $h_{ff}(b)$ (Section 4.3). We also experimented with h_{max} and h_{add} for the two cases (state-based and belief-based), but both performed much worse than h_{ff} . To avoid overcrowding the table below we do not report these experiments. We also use a flat uninformative heuristic, that gives a cost of 1 to all non-goal beliefs.

Procedure: For each heuristic method we run RTDP-BEL for at most 300 seconds of computation time. We stop the computation repeatedly to evaluate the current policy, to see if the algorithm has converged. The time required for policy evaluation is not considered a part of the runtime of the algorithm. Policy evaluation is run for each possible start state, or for 100 iterations, if there are less than 100 initial states. Each iteration runs from a start state until the current belief satisfies the goal. We stop an iteration if it exceeded 500 steps, assuming a loop. We select the action a with the best $C(a) + \sum_o pr(o|b, a)V(b_o^o)$, given the current value function V . During policy evaluation (Eqn. 3), if a belief b is encountered for which $V(b)$ was not yet computed, we use the value that the evaluated heuristic assigns to that belief instead. If 5 consecutive policy evaluations succeeded reaching the goal in all iterations, and the average cost of the policy evaluation has not changed by more than 1%, we assume that RTDP-BEL has converged. We then run a policy evaluation over the final value function for 1000 iterations, and compute the average cost, and the number of failed iterations. We report averages over 50 executions.

Domains: While there exist a set of standard benchmarks for evaluating POMDP algorithms, such as RockSample, or LaserTag [28], these benchmarks do not exhibit significant effort for information gathering. As such domains are the focus of our attention in this paper, we use a set of domains adapted from the contingent planning literature [1] that exhibit these properties.

The Wumpus domain (denoted W in the tables below) is an interesting domain where lengthy information gathering sequences of actions are needed, together with multiple sensing actions are different positions. While for each state there is a simple path to the goal, many observations are needed to understand which path applies for the current state. As such, it is a good testing bed for the ability of

¹ <https://github.com/aiplan4eu/up-cpor>, upon acceptance.

Table 1: Problem properties. D/S denotes deterministic / stochastic transitions. U/N denotes uniform / non-uniform initial belief. L denotes localize, EL denotes logistics, CB colorballs, B blocks, D doors, M maze, and W denotes wumpus. $|B|$ estimates the number of beliefs observed during the execution of an optimal policy.

Name	Type	$ P $	$ S $	$ A $	$ B $	$E[C]$
L5	DU	21	70	9	58	12.8
L5	SU	21	71	9	69	20.01
L7	DU	34	119	9	102	18.44
L7	SU	34	120	9	131	22.65
L9	DU	46	180	9	164	23.97
L9	SU	46	182	9	190	38.87
EL	DU	19	24,000	96	1097	20.52
CB	DU	14	4240	56	244	15.27
B7	DU	63	37,477	504	57	8.57
B7	DN	63	37,477	504	154	22.65
B7	SU	63	37,477	504	186	8.31
D7	DU	70	9457	336	2532	30.89
D7	DN	70	9457	336	2597	29.27
D7	SN	76	8×10^8	426	19	20.45
M7.1	DU	61	2368	308	1526	30.2
M7.2	DU	63	4736	308	3046	32.05
W4	DU	38	540	81	82	12.33
W4	DN	38	540	81	94	12.22
W5	DU	47	4968	131	259	16.94
W5	DN	47	4968	131	318	14.71
W6	DU	65	33,264	191	845	21.51
W6	DN	65	33,264	191	993	20.21

heuristics to consider sensing actions. A parameter n defines the size of the grid and the number of possible Wumpi, and hence, the amount of needed observations. We create two versions — one where there is a uniform distribution of possible states, and one where the Wumpi are more likely to be on one side. Actions have deterministic effects.

Although Wumpus requires much information gathering, it does not involve value of information, because sensing actions are mandatory. We hence create a maze domain, where the agent must reach a goal position. In the maze there are several bottlenecks where the agent must pass through one of two cells, one easy to traverse (success probability of 1.0) and the other not. If the cell is hard, or if the agent does not know whether the cell is easy, it can still move through it, but the success probability drops to 0.1. For each possible state there is a path that consists only of easy cells. There are a few cells from which the agent can observe to see whether some cells are easy. The agent must balance between the cost of visiting these observation cells, and moving towards the goal.

Localize (denoted L) is a domain where the agent must navigate a simple maze, with uncertainty about its position, observing nearby walls. We add non uniform stochastic transitions, so the agent should consider the expected cost to the goal. That is, some paths have a higher probability of success than others. Localize domains are similar to the well known Hallway problems [20] which are still relatively challenging for modern point-based solvers due to high uncertainty.

In the Colorballs (denoted CB) domain the agent must search colored balls in a small grid and move them to designated positions. The Elog domain (denoted EL) is a standard logistics problem. In these two domains we have uniform initial states and deterministic actions.

Blocks (denoted B) is a standard Blocksworld problem, with uncertainty about the positions of blocks, and stochastic outcomes where blocks may fall on the table. In Doors an agent must move through a grid, establishing which doors are open. We also add a version where doors can be opened. Some doors are easier to open and some more difficult, implemented using stochastic effects. For some doors the agent cannot know whether they are easy or not to open, but can deduce it from succeeding or failing to open a closed door.

Table 1 shows the properties of the domains. There are several types of some problems, exhibiting either deterministic or stochastic transitions (denoted D or S under Type), and either uniform or non-uniform initial belief (denoted U or N). For problem size we report the number of facts ($|P|$) in the problem description, without facts

Table 2: Comparing the time (seconds) until convergence of delete relaxation heuristics. For both state-based and belief-based we compare h_{max} , h_{add} , h_{ff} .

Problem		Belief-based			State-based		
Name	Type	Max	Add	FF	Max	Add	FF
L5	DU	0.42	0.43	0.22	0.03	0.03	0.03
L5	SU	1.44	1.47	1.23	0.2	0.18	0.29
L7	DU	3.04	3.08	1.38	0.11	0.12	0.12
L7	SU	22.8	21.6	8.3	2.89	2.54	9.27
L9	DU	16.1	15.3	10.1	0.28	0.27	0.29
L9	SU	74.4	72.9	56.4	32.1	31.8	64.4
EL	DU	15.7	8.07	11.8	14	6.88	12.4
CB	DU	1.72	1.18	1.27	1.72	1.65	1.69
B7	DU	0.97	0.24	0.04	1.18	0.42	0.27
B7	DN	1.15	0.33	0.03	1.25	0.49	0.3
B7	SU	231	288.9	24.2	193.6	291.8	210.3
D7	DU	22.4	20.7	15.8	9.84	9.88	9.5
D7	DN	8.27	8.38	4.2	5.69	5.38	5.48
D7	SN	0.05	0.05	0.02	0.74	0.75	0.74
M7.1	DU	23.4	22.2	2.44	8.05	8.2	7.99
M7.2	DU	99.1	62.5	3.79	27.7	31	31.1
W4	DU	0.22	0.22	0.16	0.21	0.2	0.2
W4	DN	0.16	0.15	0.07	0.21	0.2	0.2
W5	DU	6.8	7.05	4.92	8.57	8.58	7.83
W5	DN	6.4	6.64	3.3	6.02	6.13	6.16
W6	DU	152.4	127.1	37.2	150.7	133.2	133
W6	DN	146.7	124.3	25.2	135.3	142.3	131.7

Table 3: Comparing the time (seconds) of h_{ff} to MDP heuristics.

Name	Type	$h_{ff}(b)$	$h_{ff}(s)$	Q_{MDP}	ML	Flat
L5	DU	0.22	0.03	0.01	0.003	0.02
L5	SU	1.23	0.28	0.03	0.02	0.27
L7	DU	1.38	0.12	0.05	0.01	0.11
L7	SU	8.3	13.5	0.23	0.15	13.5
L9	DU	10.1	0.32	0.12	0.01	0.29
L9	SU	83.3	89.2	0.19	0.27	114.3
EL	DU	12.1	17	6.02	5.56	18.1
CB	DU	1.25	1.47	1.98	1.45	1.0
B7	DU	0.04	0.26	1.51	0.73	7.18
B7	DN	0.03	0.31	1.78	1.42	7.53
B7	SU	24.2	210.3	×	×	×
D7	DU	15.8	9.95	30.9	16.9	27.3
D7	DN	4.2	5.32	25.4	9.71	28.2
D7	SN	0.02	0.92	12.3	0.4	×
M7.1	DU	2.44	8.24	10.3	8.08	23.8
M7.2	DU	3.79	36.1	53.7	35.2	82.6
W4	DU	0.16	0.22	0.3	0.15	0.45
W4	DN	0.07	0.21	0.27	0.17	0.43
W5	DU	4.92	7.26	9.17	5.16	57.6
W5	DN	3.3	5.22	7.19	4.44	51.6
W6	DU	37.2	260.7	246.9	128.5	×
W6	DN	25.2	239.1	219.5	129.8	×

whose value is fixed, such as the representation of adjacent cells in a grid. Potentially, there can be $2^{|P|}$ possible states, and a naive representation of the belief state would be unmanageable. However, but due to dependencies between the facts the actual number of states is much lower. For example, in Blocks, block A cannot be at the same time on block B and on block C . Also, we can never reach a state where block A is on block B and block B is on block A . Hence, we explicitly compute the number of possible states ($|S|$) in each of the problems. As can be seen, this number is substantially lower, allowing us to represent the belief state efficiently.

The number of actions ($|A|$) may be misleading, because while in some domains, such as Blocks, there are many actions, typically only for a few of them the precondition hold at a given belief state. $|B|$ is the amount of belief states that were visited by any of the methods while acting under the final policy. RTDP-BEL must visit at least these beliefs, and their neighbors, before converging to an optimal policy. $E[C]$ is the average cost of the best policy. Methods that did not reach this expected cost within the time limit were considered as fail to converge.

5.1 Results

Table 2 shows the average runtime until convergence for our delete relaxation heuristics. For both state-based (Section 4.1) and belief-based (Section 4.3) we compare the 3 alternatives — h_{max} , h_{add} , h_{ff} . For Localize (L) problems, where all information gathering is my-

Table 4: Comparing the number of RTDP iterations until convergence of h_{ff} to MDP heuristics.

Name	Type	$h_{ff}(b)$	$h_{ff}(s)$	Q_{MDP}	ML	Flat
L5	DU	106.2	141.2	141.8	44	136
L5	SU	384	770	162.6	135.2	886
L7	DU	188	244.6	208.8	52.8	250.6
L7	SU	1942	18780	564	644.2	23560
L9	DU	287.2	418	256.2	56.4	436
L9	SU	15800	91400	401	977	137800
EL	DU	4610	7370	2800	2776	7990
CB	DU	485	1662	1624	1554	2224
B7	DU	12.0	120.6	99.2	60.8	3999
B7	DN	17.4	229.4	216.4	151.8	9308
B7	SU	4782	×	×	×	×
D7	DU	3456	6640	5948	4612	18060
D7	DN	1109	2904	2792	1672	45180
D7	SN	8.8	38.8	78.4	54.4	×
M7,1	DU	141.4	2854	1818	2004	6440
M7,2	DU	234.6	12460	10680	14560	13520
W4	DU	45.6	220.8	201.4	136.8	706
W4	DN	93.6	215.4	179.6	190.6	575
W5	DU	600	3406	3416	2262	11920
W5	DN	858	2140	2100	1868	18640
W6	DU	2012	10830	10910	8280	×
W6	DN	2326	32020	29080	25040	×

opic, the state-based methods are best. State-based FF is less successful, especially in the stochastic case. For the rest of the problems, the belief-based FF heuristic using our belief-based method is better on many domains, but truly excels on the more difficult problems — the hardest version of Blocks, with stochastic transitions, the Maze problems, that require much information gathering, and Wumpus. As such, we continue our empirical estimation focusing on h_{ff} .

Table 3 compares the runtime of h_{ff} to MDP based methods and a flat heuristic, that assigns the same heuristic value to all states. Table 4 shows the number of RTDP iterations until convergence.

Localize (L) problems have relatively small state and action spaces. The agent is unclear about its current state, but can gather information without investing much cost. Observations are necessary to achieve preconditions of actions, such as knowing that there is no wall before moving in a particular direction, and hence must be activated in the RTDP-BEL trajectories. Thus, in these problems the MDP-based techniques excel. The most likely state heuristic is best in the deterministic version, attesting to the relative simplicity of these POMDPs. Logistics (EL) has similar properties.

We report the number of iterations as there is typically a tradeoff between the ability to produce accurate heuristic estimates, and the time it takes to compute these estimations. As can be seen, in localize, $h_{ff}(b)$ often requires less iterations, but it is still much slower, because in this domain the MDP-based methods only need to take the values from the MDP value function, requiring almost no computation time.

In color blocks (CB), one must use many sensing actions, but always myopically, with no need to invest any effort in gathering information. In this domain, the flat heuristic is best, because any heuristic effort is wasted.

The blocks domain (B) is more difficult. It has a huge number of actions, far more than any other domain, and also a large state space. In this domain, the $h_{ff}(b)$ and $h_{ff}(s)$ methods, that allow RTDP-BEL to avoid exploring many irrelevant actions, work better than the MDP-based methods. For the stochastic version, only $h_{ff}(b)$ converged in time.

In doors (D) the MDP-based techniques are inferior. The stochastic version of doors has the largest state space, but the vast majority of states need not be explored, and indeed, all heuristics (Flat excluded) manage to avoid this needless exploration. Here, $h_{ff}(b)$ is fastest.

The true value of our belief-based method is revealed in domains that require significant information gathering effort — maze (M) and wumpus (W). In these domains, $h_{ff}(b)$, that manages to identify the need for information gathering rapidly, is an order of magnitude bet-

ter than all other methods, in the larger instances. The advantage is even more pronounced when observing the amount of needed iterations until convergence. In domains such as maze, methods that do not reason about the need for information realize the value of sensing only after investigating all other options. In the larger maze problem, e.g., all methods require more than 10,000 iterations before they explore the information gathering sequences, while $h_{ff}(b)$ converges after less than 150 iterations.

We also tried to run a popular point-based method, SARSOP [23] on these problems. As SARSOP does not support preconditions, we penalize executing an action is a state where the preconditions do not hold by -10,000. SARSOP also implicitly assumes that the goal state is observable. We hence add a “declare-goal” action that moves the agent to a terminal state, with a -10,000 penalty if called from a non-goal state. Unfortunately, SARSOP was unable to even load the larger problems here. For localize 9 (L9), SARSOP took 7.1 seconds to converge to the same average cost. For doors 7 (D7), SARSOP was unable to load the problem, and on the smaller doors 5 problem, SARSOP needed 6.94 seconds to converge. On wumpus 5 (W5) SARSOP took 12.11 minutes to converge, and on maze (M7,1) it needed 2.6 hours (not including the 2 hours it took to load the model).

This is not a completely fair comparison. It might be that implementing a point-based algorithm using the structured representation that we use, with action preconditions, may be competitive with RTDP-BEL. We leave such investigations for future research.

To summarize, in domains where sensing is immediate and necessary, although our $h_{ff}(b)$ method is useful, it does not present a significant advantage, if any, over the standard MDP-based methods. However, in domains where information gathering is critical, which are arguably most appropriate for a POMDP formalization, our method provides an order of magnitude improvement over MDP-based methods, that are incapable of assessing the value of information generated by long sequences of actions.

In most complex domains our method is better in terms of the amount of required iterations, while the MDP-based methods require less iterations on the simpler Localize problems. While the number of iterations is often less important in model-based POMDP applications, it can be very important in reinforcement learning (RL) applications. In such applications one must often interact with the environment, or a costly simulator, to run trajectories in the state space. Reducing the amount of such interactions (or samples), is often more important than reducing the internal computation time. This is known as the sample complexity of an algorithm [22, e.g.]. It might be that our heuristics can be even more important in RL, in cases where a model can be specified, as we require, but the transition function is unknown. We leave this to future research.

6 Conclusion

In this paper we suggested a new heuristic for estimating the value of a belief state in a POMDP, formalized as stochastic partially observable contingent planning. Our heuristic is able to estimate the cost of lengthy information gathering action sequences, and thus can direct the planner to consider such sequences. While this computation requires more time than standard heuristics, such as the popular Q_{MDP} heuristic, we show that in tasks that require significant effort for information gathering, it is worthwhile.

For future research we will investigate using our method in other scenarios, such as in point-based POMDP methods, as well as in the context of reinforcement learning.

Acknowledgments

This paper is partially supported by the ISF fund, grant no. 964/22, and by the Helmsley Charitable Trust through the ABC fund.

References

- [1] A. Albore, H. Palacios, and H. Geffner. A translation-based approach to contingent planning. In *IJCAI*, volume 9, pages 1623–1628, 2009.
- [2] A. G. Barto, S. J. Bradtke, and S. P. Singh. Learning to act using real-time dynamic programming. *Artificial intelligence*, 72(1-2):81–138, 1995.
- [3] R. Bellman. A markovian decision process. *Journal of mathematics and mechanics*, pages 679–684, 1957.
- [4] P. Bertoli, A. Cimatti, et al. Improving heuristics for planning as search in belief space. In *AIPS*, pages 143–152, 2002.
- [5] D. Bertsekas. *Dynamic programming and optimal control*, volume 4. Athena scientific, 2012.
- [6] A. Blum and M. L. Furst. Fast planning through planning graph analysis. *Artif. Intell.*, 90(1-2):281–300, 1997. doi: 10.1016/S0004-3702(96)00047-1. URL [https://doi.org/10.1016/S0004-3702\(96\)00047-1](https://doi.org/10.1016/S0004-3702(96)00047-1).
- [7] O. Blumenthal and G. Shani. Rollout heuristics for online stochastic contingent planning. In *Workshop on Agents and Robots for reliable Engineered Autonomy*, 2023.
- [8] O. Blumenthal and G. Shani. Domain independent heuristics for online stochastic contingent planning. *Annals of Mathematics and Artificial Intelligence*, 2024.
- [9] B. Bonet and H. Geffner. Planning with incomplete information as heuristic search in belief space. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems*, pages 52–61, 2000.
- [10] B. Bonet and H. Geffner. Planning as heuristic search. *Artificial Intelligence*, 129(1-2):5–33, 2001.
- [11] B. Bonet and H. Geffner. Solving pomdps: Rtdp-bel vs. point-based algorithms. In *IJCAI*, pages 1641–1646. Pasadena CA, 2009.
- [12] B. Bonet and H. Geffner. Factored probabilistic belief tracking. In S. Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 3045–3052, 2016.
- [13] R. I. Brafman and G. Shani. Online belief tracking using regression for contingent planning. *Artificial Intelligence*, 241:131–152, 2016.
- [14] D. Bryce, S. Kambhampati, and D. E. Smith. Planning in belief space with a labelled uncertainty graph. In *AAAI Workshop on Learning and Planning in Markov Decision Processes*, 2004.
- [15] D. Bryce, S. Kambhampati, and D. E. Smith. Planning graph heuristics for belief space search. *Journal of Artificial Intelligence Research*, 26: 35–99, 2006.
- [16] H. Geffner and B. Bonet. Solving large pomdps using real time dynamic programming. In *Working Notes Fall AAAI Symposium on POMDPs*, volume 218, 1998.
- [17] P. Haslum, N. Lipovetzky, D. Magazzeni, C. Muise, R. Brachman, F. Rossi, and P. Stone. *An introduction to the planning domain definition language*, volume 13. Springer, 2019.
- [18] M. Hauskrecht. Value-function approximations for partially observable markov decision processes. *Journal of artificial intelligence research*, 13:33–94, 2000.
- [19] J. Hoffmann and B. Nebel. The ff planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14: 253–302, 2001.
- [20] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- [21] S.-K. Kim, O. Salzman, and M. Likhachev. Pomhdp: Search-based belief space planning using multiple heuristics. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 734–744, 2019.
- [22] H. Kumar, A. Koppel, and A. Ribeiro. On the sample complexity of actor-critic method for reinforcement learning with function approximation. *Machine Learning*, pages 1–35, 2023.
- [23] H. Kurniawati, D. Hsu, and W. S. Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, volume 2008. Citeseer, 2008.
- [24] J. Pineau, G. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, 27: 335–380, 2006.
- [25] P. Poupart, K.-E. Kim, and D. Kim. Closing the gap: Improved bounds on optimal pomdp solutions. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 21, pages 194–201, 2011.
- [26] S. Richter, M. Helmert, and M. Westphal. Landmarks revisited. In *AAAI*, volume 8, pages 975–982, 2008.
- [27] J. C. Saborío and J. Hertzberg. Efficient planning under uncertainty with incremental refinement. In *Uncertainty in Artificial Intelligence*, pages 303–312. PMLR, 2020.
- [28] G. Shani, J. Pineau, and R. Kaplow. A survey of point-based pomdp solvers. *Autonomous Agents and Multi-Agent Systems*, 27:1–51, 2013.
- [29] D. Silver and J. Veness. Monte-carlo planning in large pomdps. *Advances in neural information processing systems*, 23, 2010.
- [30] T. Smith and R. G. Simmons. Heuristic search value iteration for pomdps. In D. M. Chickering and J. Y. Halpern, editors, *UAI '04, Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence, Banff, Canada, July 7-11, 2004*, pages 520–527, 2004.
- [31] A. Somani, N. Ye, D. Hsu, and W. S. Lee. Despot: Online pomdp planning with regularization. *Advances in neural information processing systems*, 26, 2013.
- [32] E. J. Sondik. *The optimal control of partially observable Markov processes*. Stanford University, 1971.
- [33] M. T. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for pomdps. *Journal of artificial intelligence research*, 24: 195–220, 2005.
- [34] S. To, T. Son, and E. Pontelli. On the effectiveness of belief state representation in contingent planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, pages 1818–1819, 2011.
- [35] M. Wang and R. Dearden. Planning with state uncertainty via contingency planning and execution monitoring. In M. R. Genesereth and P. Z. Revesz, editors, *Proceedings of the Ninth Symposium on Abstraction, Reformulation, and Approximation, SARA 2011, Parador de Cardona, Cardona, Catalonia, Spain, July 17-18, 2011*, 2011.