



Optimal sensing via multi-armed bandit relaxations in mixed observability domains

Citation

Lauri, M., & Ritala, R. (2015). Optimal sensing via multi-armed bandit relaxations in mixed observability domains. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 26-30 May 2015, Seattle, WA (Vol. 2015-June, pp. 4807-4812) <https://doi.org/10.1109/ICRA.2015.7139867>

Year

2015

Version

Peer reviewed version (post-print)

Link to publication

[TUTCRIS Portal \(http://www.tut.fi/tutcris\)](http://www.tut.fi/tutcris)

Published in

2015 IEEE International Conference on Robotics and Automation (ICRA), 26-30 May 2015, Seattle, WA

DOI

[10.1109/ICRA.2015.7139867](https://doi.org/10.1109/ICRA.2015.7139867)

Copyright

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Take down policy

If you believe that this document breaches copyright, please contact cris.tau@tuni.fi, and we will remove access to the work immediately and investigate your claim.

Optimal Sensing via Multi-armed Bandit Relaxations in Mixed Observability Domains

Mikko Lauri and Risto Ritala

Abstract—Sequential decision making under uncertainty is studied in a mixed observability domain. The goal is to maximize the amount of information obtained on a partially observable stochastic process under constraints imposed by a fully observable internal state. An upper bound for the optimal value function is derived by relaxing constraints. We identify conditions under which the relaxed problem is a multi-armed bandit whose optimal policy is easily computable. The upper bound is applied to prune the search space in the original problem, and the effect on solution quality is assessed via simulation experiments. Empirical results show effective pruning of the search space in a target monitoring domain.

I. INTRODUCTION

Deploying autonomous agents such as robots equipped with an appropriate set of sensors allows automated execution of various information gathering tasks. The tasks can include monitoring and identification of spatio-temporal processes, automated exploration, or other data collection campaigns in environments where human presence is undesired or infeasible. Robots are mobile sensor platforms whose actions are optimized to maximize the informativeness of measurement data.

As the target state is not known, a probability density function (pdf) over the state, called a belief state, is maintained. Information conveyed by measurement data is incorporated into the belief state by Bayesian filtering. Assuming Markovian dynamics and conditional independence of measurement data given the system state, the problem is a partially observable Markov decision process, or POMDP [1].

Optimal information gathering has been studied in the context of sensor management [2], and a review of applying POMDPs for sensor management is presented in [3]. The problem is formulated as a decision process under uncertainty. The goal is to find a control policy mapping belief states to actions, that when followed maximizes the expected sum of discounted rewards over a horizon of time. The reward associated with an action may depend either on the true state of the system or the belief state. The former can encode objectives such as reaching a favorable state or avoiding costly ones, useful e.g. for navigation and obstacle avoidance. The latter option allows information theoretic rewards, such as mutual information, applied in various sequential information gathering problems in robotics, see e.g. [4], [5], [6]. Indefinite-horizon problems that terminate

when a special stopping action is executed are a natural model for tasks that may be stopped once a certain level of confidence about the state is reached [7].

Finding optimal policies for POMDPs is computationally hard [8], and several approximate methods have been suggested. Point-based algorithms [9], [10] track so-called alpha vectors at a set of points in the belief space. The alpha vectors may then be used to approximate the optimal policy at any belief state. Online planning methods [11] find an optimal action for the current belief state instead of a representation of the optimal policy. The problem is cast as a search over the tree of belief states reachable from the current belief state under various action-observation histories. Combining online methods with Monte Carlo simulations to evaluate utility of actions has led to approximate algorithms able to handle problems with up to 10^{52} states [12].

In mixed observability domains a part of the state space is fully observable. The belief space is a union of low-dimensional subspaces, one for each value of a fully observable state variable. Robotic systems often exhibit mixed observability which may be exploited to derive efficient POMDP algorithms [13].

A multi-armed bandit (MAB) is a model for sequential decision-making also applied in sensor management [2]. A decision-maker plays one arm of the MAB and collects a reward depending on the state of the arm. The arm then randomly transitions to a new state while other arms remain stationary. Solutions to MABs are index policies that are easier to compute than solutions to general POMDPs [14].

Most of the aforementioned research applies reward functions that only depend on the true state and action. The expectation of the reward is linear in the belief state, a feature leveraged by many of the solution algorithms. Information theoretic quantities such as entropy and mutual information that would be useful as reward functions in optimal sensing problems are nonlinear in the belief state. Classical POMDP algorithms cannot be applied to solve such problems.

In this paper, we study POMDPs with mixed observability with mutual information as the reward function. As such, our approach is especially suited for optimal sensing problems in robotics domains. We remove constraints on available actions to obtain a relaxed problem. The optimal value of the relaxed problem obtained is an upper bound on the optimal value in the POMDP. We identify the conditions under which the relaxed problem is a MAB and has an easily computable optimal solution. The upper bound is applied in an online planning algorithm to prune the search space.

The paper is organized as follows. In Section II, the mixed-

M. Lauri and R. Ritala are with Department of Automation Science and Engineering, Tampere University of Technology, P.O. Box 692, FI-33101 Tampere, Finland. Email: mikko.lauri@tut.fi, risto.ritala@tut.fi

observability POMDP is defined. In Section III, methods for solving the problem are discussed. In Section IV, two relaxations are derived that provide upper bounds for the optimal value function. Section V determines the conditions under which the relaxations are MABs. Empirical results are provided in Section VI. Section VII concludes the paper.

II. A MIXED OBSERVABILITY POMDP

Notation. We denote random variables and sets by uppercase letters, and realizations of random variables and members of sets as lowercase letters. Time instants are distinguished by writing e.g. x and x' for realizations at time t and $(t + 1)$, respectively.

An agent, e.g. a robot or another sensor platform, has an internal state $x \in X$ that captures the dynamics and constraints of operating on-board sensors and other devices. The internal state evolves according to a deterministic dynamics model D_X , defined $x' = f_a(x)$ where $a \in A(x)$ is a control action in the finite set of actions allowed in internal state x .

Let $Y = \{Y_1, Y_2, \dots, Y_n\}$, $X \cap Y = \emptyset$, denote a set of random inference variables an agent wishes to obtain information about. The dynamics of the variables are governed by a stochastic model D_Y , defined as a Markov chain $p(y' | y, a)$. The complete state of the system is $s \in S = X \times Y$.

The problem features mixed observability, where the internal state is fully observable and the inference variables are partially observable. The agent's observations $z' \in Z$ follow an observation model O , defined by $p(z' | y', a)$.

The agent maintains a belief state $b = (x, p(y)) \in B$, consisting of the deterministic, fully observable internal state and a pdf over Y . The initial belief state b_0 is given. Given a belief state b , an action a , and an observation z' , the belief state at the next time instant is given by the belief update equation $b' = \tau(b, a, z') = (x', p(y' | z', a, b))$ where $x' = f_a(x)$, and the pdf over the inference variables is obtained from a Bayesian filter

$$p(y' | z', a, b) = \frac{1}{\eta} p(z' | y', a) p(y' | a, b) \quad (1)$$

where $p(y' | a, b) = \sum_{y \in Y} p(y' | y, a) p(y)$ is the predictive pdf and $\eta = p(z' | a, b)$ is the normalization factor denoting the prior probability of observing z' . Given any sequence of actions and observations, there is no uncertainty about the resulting internal state x . Thus we can equivalently define the set of allowed actions $A(x)$ via the belief state as A_b .

The agent's objective is encoded by a reward function R . The objective is to maximize the expected sum of discounted rewards over a horizon of T decisions. The discount factor is $\gamma \in [0, 1]$.

We consider belief-dependent reward functions. Let $R(b, a) = I(Y; Z | a)$, i.e. the mutual information (MI) between the posterior state and observation. MI is defined

$$I(Y; Z | a) = H(Y | a) - \mathbb{E}_Z [H(Y | z', a)], \quad (2)$$

where $H(Y | a)$ is the entropy of the predictive pdf $p(y' | a, b)$ and the second term is the expected entropy of the posterior pdf (1) under the prior pdf $p(z' | a, b)$.

The problem $P = \langle S, A, Z, D, O, R, b_0, \gamma \rangle$ where $D = D_X \times D_Y$ is an instance of a POMDP. By Bellman's principle of optimality [15] the solution may be found via a backward in time recursion procedure known as value iteration. An optimal value function $V_t^* : B \rightarrow \mathbb{R}$ maps a belief state to its maximum expected sum of discounted rewards when an optimal policy is followed for the next t decisions. Optimal value functions are computed by

$$Q_t(b, a) = R(b, a) + \gamma \sum_{z \in Z} p(z | a, b) V_{t-1}^*(b') dz \quad (3a)$$

$$V_t^*(b) = \max_{a \in A_b} Q_t(b, a), \quad (3b)$$

starting from $Q_1(b, a) = R(b, a)$. The optimal policy $\pi_t^* : B \rightarrow A_b$ for t remaining decisions is found by extracting the argument a maximizing $Q_t(b, a)$. The recursion is continued up to V_T .

III. SOLVING POMDPs WITH BELIEF-DEPENDENT REWARDS

In most POMDPs, the reward function is state-dependent and its expectation is linear in the belief state. The finite-horizon optimal value function then has a finite representation by a convex hull of a set of hyperplanes over the belief space [16]. Many exact [17] and approximate [18], [10], [9] offline algorithms for POMDPs rely on this piecewise linearity and convexity of the value function. Reward functions such as mutual information and entropy that are useful in optimal sensing problems are nonlinear in the belief state. Thus, these offline algorithms are not applicable to solve the recursion (3) with a belief-dependent reward function.

Online planning methods [11] find an optimal action for the current belief state instead of a closed form representation of the optimal policy. As explicit representations of policies are not required, a nonlinear belief-dependent reward function does not constitute any additional difficulty.

In online planning, a tree graph of belief states reachable from the current belief state is constructed. The current belief state is the root of the tree, and belief states computed via $\tau(b, a, z')$ are added as child nodes of node b . When a desired search depth is reached, the values from the leaves of the tree are propagated back to the root according to (3).

Suboptimal actions may sometimes be pruned from the search tree by branch-and-bound pruning when the optimal value for executing action a in belief state b , $Q_t(b, a)$, has an upper bound $U(b, a)$ and a lower bound $L(b, a)$. For a given a and any $\hat{a} \neq a$, if $U(b, a) \leq L(b, \hat{a})$ then action a is suboptimal at b and all its successor nodes may be pruned from the tree. The bounds may similarly be propagated via (3). The number of belief states in the search tree is reduced.

Alternatives to online tree search include e.g. specialized approximate methods [19], however limited to small problems, open-loop approximation applied with the receding horizon control principle [6], or reduced value iteration [5] for Gaussian beliefs over Y in a mixed-observability case. For a theoretical treatment of nonlinear but convex reward functions in POMDPs, we refer the reader to [20].

IV. BOUNDS FOR THE VALUE FUNCTION

The optimal policy π_t^* attains the optimal value for all belief states. Then any other policy π_t achieves a value that is a lower bound on the optimal value. A simple choice is to set π_t as the greedy one-step look-ahead policy $\pi_G(b) = \operatorname{argmax}_{a \in A_b} R(b, a)$. Other options include random policies or blind policies [18] always executing a single fixed action.

Upper bounds are found by deriving two relaxed versions of the original POMDP problem by removing constraints on the applicable actions. The set of internal states reachable from a subset $X_S \subseteq X$ in a single time step is

$$F(X_S) = \bigcup_{\forall x \in X_S} \bigcup_{\forall a \in A(x)} f_a(x). \quad (4)$$

The set of internal states reachable in k steps from X_S is

$$F^k(X_S) = \underbrace{F \circ \dots \circ F}_{k \text{ times}}(X_S). \quad (5)$$

The first relaxation is obtained by removing all constraints imposed by the internal state as follows.

Universal sensor relaxation. Given a POMDP problem $P = \langle S, A, Z, D, O, R, b, \gamma \rangle$, its universal sensor relaxation is $P_u = \langle Y, \hat{A}, Z, D_Y, O, R, p(y), \gamma \rangle$, where $\hat{A} = \bigcup_{x \in X} A(x)$ contains all actions, D_Y is the stochastic part from D , and $b = (x, p(y))$ is replaced by $p(y)$.

When we consider only actions applicable in the internal states reachable within $k = T - t$ decisions, where t is the current time step, we obtain the k -step sensor relaxation.

k -step sensor relaxation. Given a POMDP problem $P = \langle S, A, Z, D, O, R, b, \gamma \rangle$, its k -step sensor relaxation is $P_k = \langle Y, \hat{A}_k, Z, D_Y, O, R, p(y), \gamma \rangle$, where $\hat{A}_k = \bigcup_{i \in F^k(\{x\})} A(i)$ is the set of all actions possible in the internal states reachable within k time steps from the current internal state x , and D_Y and $p(y)$ are as for P_u .

As $A_b \subseteq \hat{A}_k \subseteq \hat{A}$, the optimal value in either relaxed problem is greater than or equal to the optimal value in the original problem. Let V^* , V_u^* , and V_k^* denote the optimal value functions for P , P_u , and P_k , respectively, and let V^{π_G} denote the value function for the greedy policy in P for a given $1 \leq t \leq T$. Now

$$V^{\pi_G}(b) \leq V^*(b) \leq V_k^*(b) \leq V_u^*(b) \quad \forall b \in B \quad (6)$$

holds for the optimal value and the bounds.

V. MULTI-ARMED BANDIT INDEX POLICIES FOR POMDP RELAXATIONS

Both relaxations defined above are POMDPs themselves. Solving even the relaxed problems may thus be a computationally intractable task. This motivates identifying POMDPs whose relaxations have easily computable optimal policies.

In a multi-armed bandit (MAB) problem, a decision-maker plays one arm of the MAB and collects a reward depending on the state of the arm. Four requirements distinguish MABs among general stochastic control problems [2]: 1) exactly

one machine is played by the agent per action, and the state of that machine evolves such that the agent may not affect it, 2) machines not played remain in their current state, 3) the machines are independent, and 4) the machines that are not played do not contribute any reward. Gittins [21] showed that the optimal policies in MABs are so-called greedy index allocation policies. For each arm, an allocation index known as the Gittins index is calculated with the optimal selection yielding the highest index value. Index policies are optimal when actions are not irrevocable [14], [2]: any action is available at any stage, and may be chosen at a later stage with the same reward, excluding the effect of the discount factor. Index policies are usually much easier to compute than backward induction solutions of POMDPs [14].

An index policy is in general not optimal for the mixed-observability POMDP of Section II, as actions are irrevocable due to the constraints imposed by the internal state. However, both of the relaxations P_u and P_k have a fixed action space. The following three properties are required for the relaxations to be MABs. Results are derived for P_u , and they hold for the more restricted case P_k as well.

Property 1. Each $a \in \hat{A}$ is related to $Y_a \subset Y$, such that $Y_a \neq \emptyset$ and $i, j \in \hat{A} : i \neq j \Rightarrow Y_i \cap Y_j = \emptyset$.

Property 2. Given $a \in \hat{A}$, each $y'_i \in Y_a$ is conditional on the values $u_i \in U_i \subseteq Y_a$ of some subset U_i of the inference variables in Y_a , and $y_i \notin Y_a$ are stationary, i.e.

$$p(y' | y, a) = \prod_{y_i \in Y_a} p(y'_i | u_i) \prod_{y_i \in Y \setminus Y_a} \delta(y'_i - y_i), \quad (7)$$

where δ is the Dirac delta function.

Property 3. For $a \in \hat{A}$, the observation is conditional on $y'_a \in Y_a$, i.e.

$$p(z' | y', a) = p(z' | y'_a). \quad (8)$$

Corollary 1. When $p(y) = \prod_{a \in \hat{A}} p(y_a)$, $y_a \in Y_a$, i.e. the prior on inference variables is independent between the subsets Y_a and properties 2 and 3 hold, the independence is preserved in the posterior,

$$p(y' | z', a, b) = \prod_{k \in \hat{A}} p(y'_k | z', a, b). \quad (9)$$

Furthermore,

$$I(Y; Z | a) = I(Y_a; Z). \quad (10)$$

Equation (9) is seen to hold applying (1) to the given prior with models satisfying (7) and (8). Equation (10) is seen to hold through two steps. First, due to the independence structure of the prior and posterior, $H(Y | a) = \sum_{k \in \hat{A}} H(Y_k | a)$ and similarly for $H(Y | z', a)$. Second, by (1) we see from (9) that for $k \neq a : p(y'_k | z', a, b) = p(y_k) \Rightarrow H(Y_k | a) = H(Y_k | z', a)$. Applying these steps to (2) leads to (10).

We now state our main result determining the conditions under which a POMDP relaxation is a MAB.

Proposition 1 (MAB equivalence of POMDP relaxations). *When properties 1-3 are fulfilled, and the prior is $p(y) = \prod_{a \in \hat{A}} p(y_a)$, $y_a \in Y_a$, the relaxations P_u and P_k are multi-armed bandit problems.*

Proof. (Sketch). Consider the four requirements for MABs introduced above. Property 1 establishes the "arms" of the bandit, partly satisfying requirement 1. The rest of requirements 1 and 2 are satisfied by Properties 2 and 3, which establish the states of the bandit arms as $p(y_a)$, $y_a \in Y_a$, $a \in \hat{A}$. Requirement 3 is satisfied by the independence properties in the first part of Corollary 1. The latter part of the corollary shows that requirement 4 is satisfied. \square

When the proposition holds, the optimal policies for P_u and P_k are greedy index policies with values $V_u^* = V_u^{\pi_G}$ and $V_k^* = V_k^{\pi_G}$, respectively. These optimal values are thus much easier to compute than for general POMDPs.

Let us consider the following example problem.

Monitoring reactive targets. *An agent is located at $x \in X = \{1, 2, \dots, M\}$. At every time step the agent may either stay where it is or move to one of the neighboring locations $N(x) \subset X$. The applicable actions are $A(x) = x \cup N(x)$. Let $Y = \{Y_1, Y_2, \dots, Y_M\}$, with Y_i assuming value 1 if a target is present at location i and 0 if not. Each target Y_i reacts to the agent's presence such that*

$$p(y'_i | y_i, a) = \begin{cases} w_i(y'_i | y_i) & \text{if } a = i \\ r_i(y'_i | y_i) & \text{if } a \neq i \end{cases} \quad (11)$$

The agent records measurements in $Z = \{0, 1\}$ according to

$$p(z' = 0 | y', a) = \begin{cases} 1 - q_- & \text{if } y'_a = 0 \\ q_+ & \text{if } y'_a = 1 \end{cases}, \quad (12)$$

where $q_- < 0.5$, $q_+ < 0.5$ are the false negative and positive probabilities, respectively, and $p(z' = 1 | y', a) = 1 - p(z' = 0 | y', a)$. The reward function is (2).

Consider the relaxations P_u and P_k of this problem. Property 1 is immediately seen to be satisfied. Property 2 is satisfied if $r_i(y'_i | y_i) = \delta(y'_i - y_i)$ for $1 \leq i \leq M$, i.e. if targets remain stationary when the agent is not present, while $w_i(y'_i | y_i)$ may be chosen freely. Property 3 is satisfied as the observation only depends on the value y'_a .

VI. EMPIRICAL EVALUATION

We ran simulation experiments on the monitoring problem defined above. There were $|Y| = M = 36$ inference variables, arranged on a rectangular two-dimensional four-connected grid. The agent was allowed to move on this grid and sense the targets. We examined two cases. In the first case, all of the properties 1-3 were satisfied. In the second case, we relaxed Property 2 by allowing all inference variables to change state. In all cases, the optimization horizon T was varied from 1 to 6 decisions. The other parameters were $q_- = 0.05$, $q_+ = 0.05$, $\gamma = 0.95$.

We implemented the real-time belief space search (RTBSS) algorithm of [22] as presented in [11]. RTBSS

implements an online search of belief states reachable from the current belief state, and applies lower and upper bounds to prune suboptimal actions. We applied the greedy lower bound V^{π_G} (Section IV) and upper bounds $V_k^{\pi_G}$ or $V_u^{\pi_G}$ (Section V). We compared this approach to an exhaustive search of all reachable belief states equivalent to using lower and upper bounds $(-\infty, \infty)$, and to the POMCP algorithm [12], which gives a recommendation on the next action to execute based on a series of Monte Carlo (MC) simulations.

A. Case 1: Properties 1-3 satisfied

We defined $r_i = \delta(y'_i - y_i)$, and w_i were two-state Markov chains with parameters p_{01}^i, p_{11}^i , where p_{jk}^i denotes the probability that Y_i transitions from j to k . For each $1 \leq i \leq M$, we sampled uniformly at random $p_{01}^i \in [0.0, 0.2]$, $p_{11}^i \in [0.8, 1.0]$. A set of 1000 initial belief states $(x_0, p(y_0))$ satisfying the independence assumption between inference variables was sampled uniformly at random.

As Proposition 1 holds, $V_u^* = V_u^{\pi_G}$ and $V_k^* = V_k^{\pi_G}$, and the greedy MAB policies give valid upper bounds, see (6). Applying RTBSS with these bounds hence always finds the optimal solution, which was verified in our simulations. The number of visited nodes in the search tree for each of the 1000 belief states is shown in Fig. 1 for $3 \leq T \leq 6$ and both upper bounds. Since the bound V_k^* is tighter, applying it results in a lower or equal number of visited nodes than V_u^* . For comparison, the average number of visited nodes for the exhaustive search is shown in Table I. We note that applying either bound greatly reduces the number of visited nodes, in some cases by up to an order of magnitude. Although the reduction in the number of visited nodes is substantial, evaluating the bounds has a computational cost that must be balanced with the savings from visiting fewer nodes. This point is discussed in more detail in the next subsection.




Fig. 1. The number of search tree nodes expanded by RTBSS for $3 \leq T \leq 6$ with the upper bound from P_u (x -axis) or P_k (y -axis). The diagonal line shows where the two values are equal.

POMCP recommendations coincide with the optimal action more reliably when the number of MC simulations is increased and the optimization horizon T is short, see Table II. We compared the values of optimal actions to those recommended by POMCP when the two differed. The difference

TABLE I

AVERAGE NUMBER OF NODES EXPANDED BY EXHAUSTIVE SEARCH.

	$T = 3$	$T = 4$	$T = 5$	$T = 6$
Nodes	$4.0 \cdot 10^2$	$3.8 \cdot 10^3$	$3.6 \cdot 10^4$	$3.5 \cdot 10^5$

TABLE II

PERCENTAGE OF POMCP RECOMMENDATIONS AGREEING WITH OPTIMAL.

MC simulations	$T = 2$	$T = 3$	$T = 4$	$T = 5$	$T = 6$
10^1	40.7	36.8	31.3	28.1	29.2
10^2	57.4	50.1	45.3	43.2	40.1
10^3	69.5	62.2	54.0	47.6	46.8
10^4	75.0	74.9	69.1	63.5	59.1

between the two values is the performance loss, for which we computed the mean values and worst-case maximum values. The results are shown in Table III. Performance loss tends to be greater for fewer MC simulations and a greater optimization horizon T . As the number of MC simulations increases the mean performance loss is low, indicating that on average POMCP performs very well compared to the optimal solution. However, even if the mean performance loss is low, the worst case performance loss from following POMCP recommendations may be significantly greater. In problems where suboptimal actions may lead to unacceptable performance loss, methods such as RTBSS with valid bounds may be preferable to POMCP.

B. Case 2: Property 2 not satisfied

We next examined the case where Property 2 was not satisfied. We set $r_i = w_i$ for each i . Each of the dynamics models was a two-state Markov chain. We considered three subcases distinguished by the rate of the state transitions: slow, medium or fast. For slow dynamics, the parameters were sampled for each i uniformly at random such that $p_{01}^{i,slow} \in [0.0, 0.2]$, $p_{11}^{i,slow} \in [0.8, 1.0]$, for medium dynamics $p_{01}^{i,med} \in [0.2, 0.4]$, $p_{11}^{i,med} \in [0.6, 0.8]$, and for fast dynamics $p_{01}^{i,fast} \in [0.4, 0.6]$, $p_{11}^{i,fast} \in [0.4, 0.6]$. Each experiment was repeated for 1000 randomly sampled initial belief states and dynamics models. All beliefs satisfied the independence assumption between inference variables.

The problem is quite similar to the one in Subsection VI-A, and POMCP performance was also observed to be very good on average. The MAB equivalence, Proposition 1, is now not satisfied for the relaxed problems. Thus, the upper bounds are approximate, and optimality for RTBSS cannot be guaranteed. We examined the effect that this had on solutions provided by RTBSS. The results are summarized in Table IV. The table shows the percentage of solutions equal to the optimal solution in case of slow, medium or fast dynamics for either the universal sensor upper bound $V_u^{\pi_G}$ from P_u or the k -step sensor upper bound $V_k^{\pi_G}$ from P_k .

Optimal solutions are found in the majority of cases, with the percentage decreasing as the optimization horizon is

TABLE IV

PERCENTAGE OF RTBSS SOLUTIONS AGREEING WITH OPTIMAL SOLUTION WHEN PROPERTY 2 WAS NOT SATISFIED.

Dynamics	Bound	$T = 2$	$T = 3$	$T = 4$	$T = 5$	$T = 6$
Slow	P_u	100%	100%	100%	100%	100%
	P_k	100%	99.9%	99.9%	99.6%	99.6%
Medium	P_u	100%	100%	99.9%	100%	100%
	P_k	100%	99.8%	99.7%	99.1%	98.8%
Fast	P_u	100%	100%	100%	100%	100%
	P_k	100%	99.9%	99.7%	99.0%	98.9%

greater and the rate of dynamics faster. Since often $V_k^{\pi_G}(b) < V_u^{\pi_G}(b)$, it is more likely that the bound obtained from the universal sensor relaxation does not overestimate the optimal value, and consequently better agreement with the optimal solution is observed. The results suggest that it may still be reasonable to approximate upper bounds for V^* by the value of greedy policies in the relaxed problems P_u or P_k , even if their optimality cannot be guaranteed.

Efficiency of pruning the search tree was not affected significantly compared to the case of the previous subsection. Applying either bound dramatically reduced the number of visited nodes in the search tree. We examined the mean time required to find a solution for a belief state either by exhaustive search or branch-and-bound pruning. A representative comparison is presented in Fig. 2 for the case of medium dynamics. For $T < 3$, exhaustive search performs fastest: the computational burden of computing the bounds outweighs the savings from visiting fewer nodes during the search. The advantages of pruning the search tree become apparent for $T \geq 4$. At best, applying pruning is an order or magnitude faster than exhaustive search. For $T \geq 3$, the upper bound from P_k is fastest. Using the upper bound from P_u is faster than exhaustive search for $T \geq 5$. Comparing computation times between POMCP and RTBSS were not meaningful, as the experiments were run on different computer platforms with different implementations of e.g. the search trees.




Fig. 2. The mean runtime per decision in milliseconds as a function of optimization horizon T for the exhaustive and branch-and-bound search applying upper bounds from P_u or P_k .

TABLE III
PERFORMANCE LOSS OF POMCP COMPARED TO OPTIMAL.

MC simulations	$T = 2$		$T = 3$		$T = 4$		$T = 5$		$T = 6$	
	Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max
10^1	0.0749	0.3930	0.0948	0.5278	0.1061	0.5283	0.1084	0.5587	0.1175	0.6278
10^2	0.0298	0.2296	0.0518	0.3299	0.0584	0.2953	0.0642	0.3251	0.0670	0.3968
10^3	0.0122	0.1025	0.0294	0.1751	0.0402	0.2035	0.0487	0.3639	0.0518	0.2839
10^4	0.0064	0.0503	0.0168	0.0893	0.0261	0.1399	0.0317	0.1816	0.0380	0.2020

VII. CONCLUSIONS

An optimal sensing problem in a mixed-observability domain where an internal state is fully observable and a set of inference variables are partially observable was formulated as a POMDP. The objective was the sequential maximization of mutual information of the inference variables and observations. Upper bounds for the optimal value function were found by relaxing constraints of the original problem.

When three conditions are fulfilled, the relaxed problems are MABs. First, each action is related to a unique subset of inference variables. Secondly, only inference variables in the subset corresponding to the current action evolve, while the other inference variables remain stationary. Finally, observations depend only on the inference variables in the subset related to the current action. The optimal solution of a MAB problem is a greedy index allocation policy, which is much easier to find than solving a general POMDP.

The POMDP was solved by a branch-and-bound search. The effectiveness of the bounds for pruning the search space was empirically verified in a target monitoring problem. Finding an optimal action by requires searching a fraction of the reachable belief states compared to an exhaustive search. Computation time is at best an order of magnitude smaller when applying pruning. The computational savings become apparent when savings due to reduced search space size exceed the additional cost of computing the bounds.

Future work includes studying applicability of our methodology in a wider range of mixed observability domains. Motivated by positive results on optimality of greedy policies for the restless bandit problem [23], we believe there may exist more classes of stochastic control problems than currently known where a greedy policy is optimal. Identifying such classes would further expand the applicability of our results.

REFERENCES

- [1] L. Kaelbling, M. Littman, and A. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.
- [2] A. O. Hero, D. Castañón, D. Cochran, and K. Kastella, Eds., *Foundations and Applications of Sensor Management*. New York, NY: Springer, 2007.
- [3] E. K. P. Chong, C. M. Kreucher, and A. O. Hero, "Partially Observable Markov Decision Process Approximations for Adaptive Sensing," *Discrete Event Dynamic Systems*, vol. 19, no. 3, pp. 377–422, May 2009.
- [4] B. Charrow, V. Kumar, and N. Michael, "Approximate representations for multi-robot control policies that maximize mutual information," *Autonomous Robots*, vol. 37, no. 4, pp. 383–400, Aug. 2014.

- [5] N. Atanasov, J. Le Ny, K. Daniilidis, and G. Pappas, "Information Acquisition with Sensing Robots: Algorithms and Error Bounds," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, Hong Kong, China, June 2014, pp. 6447–6454.
- [6] M. Lauri and R. Ritala, "Stochastic control for maximizing mutual information in active sensing," in *ICRA 2014 Workshop on Robots in Homes and Industry: Where to Look First?*, Hong Kong, China, June 2014.
- [7] E. A. Hansen, "Indefinite-horizon POMDPs with action-based termination," in *Proceedings of the National Conference on Artificial Intelligence*, Vancouver, Canada, July 2007, pp. 1237–1242.
- [8] C. H. Papadimitriou and J. N. Tsitsiklis, "The Complexity of Markov Decision Processes," *Mathematics of Operations Research*, vol. 12, pp. 441–450, 1987.
- [9] M. T. Spaan and N. A. Vlassis, "Perseus: Randomized point-based value iteration for POMDPs," *Journal of Artificial Intelligence Research*, vol. 24, pp. 195–220, 2005.
- [10] J. Pineau, G. Gordon, and S. Thrun, "Anytime point-based approximations for large POMDPs," *Journal of Artificial Intelligence Research*, vol. 27, no. 1, pp. 335–380, 2006.
- [11] S. Ross, J. Pineau, S. Paquet, and B. Chaib-Draa, "Online planning algorithms for POMDPs," *Journal of Artificial Intelligence Research*, vol. 32, pp. 663–704, 2008.
- [12] D. Silver and J. Veness, "Monte-Carlo Planning in Large POMDPs," in *Advances in Neural Information Processing Systems 23*, Vancouver, Canada, Dec. 2010, pp. 2164–2172.
- [13] S. C. Ong, S. W. Png, D. Hsu, and W. S. Lee, "Planning under uncertainty for robotic tasks with mixed observability," *International Journal of Robotics Research*, vol. 29, no. 8, pp. 1053–1068, 2010.
- [14] A. O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3064–3075, 2011.
- [15] R. Bellman, *Dynamic Programming*. Princeton, New Jersey: Princeton University Press, 1957.
- [16] R. D. Smallwood and E. J. Sondik, "The optimal control of partially observable Markov processes over a finite horizon," *Operations Research*, vol. 21, no. 5, pp. 1071–1088, 1973.
- [17] W. S. Lovejoy, "A survey of algorithmic methods for partially observed Markov decision processes," *Annals of Operations Research*, vol. 28, no. 1, pp. 47–65, 1991.
- [18] M. Hauskrecht, "Value-function approximations for partially observable Markov decision processes," *Journal of Artificial Intelligence Research*, vol. 13, no. 1, pp. 33–94, 2000.
- [19] V. Krishnamurthy and D. V. Djonin, "Structured Threshold Policies for Dynamic Sensor Scheduling – A Partially Observed Markov Decision Process Approach," *IEEE Transactions on Signal Processing*, vol. 55, no. 10, pp. 4938–4957, Oct. 2007.
- [20] M. Araya, O. Buffet, V. Thomas, and F. Charpillat, "A POMDP Extension with Belief-dependent Rewards," in *Advances in Neural Information Processing Systems 23*, Vancouver, Canada, Dec. 2010, pp. 64–72.
- [21] J. C. Gittins, "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 148–177, 1979.
- [22] S. Paquet, B. Chaib-draa, and S. Ross, "Hybrid POMDP algorithms," in *Proceedings of The AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains*, Hakodate, Japan, May 2006, pp. 133–147.
- [23] S. H. A. Ahmad, M. Liu, T. Javidi, Q. Zhao, and B. Krishnamachari, "Optimality of myopic sensing in multichannel opportunistic access," *IEEE Transactions on Information Theory*, vol. 55, no. 9, pp. 4040–4050, 2009.