

Learning Logic Specifications for Soft Policy Guidance in POMCP

Giulio Mazzi*
University of Verona
Verona, Italy
giulio.mazzi@univr.it

Alberto Castellini
University of Verona
Verona, Italy
alberto.castellini@univr.it

Daniele Meli*
University of Verona
Verona, Italy
daniele.meli@univr.it

Alessandro Farinelli
University of Verona
Verona, Italy
alessandro.farinelli@univr.it

ABSTRACT

Partially Observable Monte Carlo Planning (POMCP) is an efficient solver for Partially Observable Markov Decision Processes (POMDPs). It allows scaling to large state spaces by computing an approximation of the optimal policy locally and online, using a Monte Carlo Tree Search based strategy. However, POMCP suffers from sparse reward function, namely, rewards achieved only when the final goal is reached, particularly in environments with large state spaces and long horizons. Recently, logic specifications have been integrated into POMCP to guide exploration and to satisfy safety requirements. However, such policy-related rules require manual definition by domain experts, especially in real-world scenarios. In this paper, we use inductive logic programming to learn logic specifications from traces of POMCP executions, i.e., sets of belief-action pairs generated by the planner. Specifically, we learn rules expressed in the paradigm of answer set programming. We then integrate them inside POMCP to provide soft policy bias toward promising actions. In the context of two benchmark scenarios, *rocksample* and *battery*, we show that the integration of learned rules from small task instances can improve performance with fewer Monte Carlo simulations and in larger task instances. We make our modified version of POMCP publicly available at https://github.com/GiuMaz/pomcp_clingo.git.

KEYWORDS

Partially Observable Markov Decision Processes, Planning Under Uncertainty, Inductive Logic Programming, Answer Set Programming, Explainable AI

ACM Reference Format:

Giulio Mazzi, Daniele Meli, Alberto Castellini, and Alessandro Farinelli. 2023. Learning Logic Specifications for Soft Policy Guidance in POMCP. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

*Authors contributed equally.

1 INTRODUCTION

Partially Observable Markov Decision Processes (POMDPs) are a popular framework for modeling systems with state uncertainty [3]. The state cannot be completely observed by the agent, therefore it is modeled as a probability distribution, called *belief*. However, computing optimal POMDP policies (exploring the full belief space) is hard [24]. Several approximate algorithms have been proposed to mitigate the computational issue, such as, Partially Observable Monte Carlo Planning (POMCP) [27] and Determinized Sparse Partially Observable Tree (DESPOT) [28]. Most existing approaches rely on exploring only a part of the belief space, e.g., with random particle sampling in POMCP; hence, a trade-off must be found between quality of policies and exploration cost. This hinders the scalability to real-world scenarios with large belief space. Furthermore, POMCP policies are black-box, hence the decision process underlying policy generation is not transparent to other agents and humans, reducing safety and trustability.

In this paper we address the aforementioned issues, introducing soft logic-based policy bias in the Monte Carlo tree exploration in POMCP. Differently from existing approaches which set bounds to the exploration process [17, 20], we use logic specifications to only suggest actions to POMCP solver, based on current belief. Furthermore, we do not require specifications to be provided by a domain expert. Instead, we propose to learn logic specifications offline from traces of executions, i.e., sequences of belief-action pairs (generated using POMCP or other solvers). We express specifications in Answer Set Programming (ASP) [18], a state-of-the-art paradigm for logic planning [8, 10, 21, 30]. We convert the belief to ASP representation, in terms of higher-level domain *features* specified by an expert user. Then, we use them to represent logic relationships between beliefs and actions. Defining features requires only basic domain knowledge (e.g., relevant domain quantities used to represent POMDP task instances). We then exploit Inductive Logic Programming (ILP) [23] to infer ASP rules from feature-action pairs.

We validate our approach on two benchmark POMDP problems, *rocksample* and *battery*. The contributions of this paper are the following: i) we learn specifications about POMDP policies in the ASP semantics from execution traces, requiring only definition of domain-relevant features from experts; ii) we use ASP statements to guide POMCP exploration and increase performance and scalability in challenging domain instances, e.g., with increased planning horizons; iii) we evaluate the impact of the quality and amount of

execution traces on learning outcome and POMCP performance. The implementation of our extended POMCP algorithm is made available.

2 RELATED WORKS

Merging MDP / POMDP solving algorithms with logic reasoning is a recent research trend [17, 20, 29]. One advantage of this approach is to guide policy search with commonsense reasoning and expert knowledge. For instance, in the REBA framework [29] ASP is used to describe spatial relations between objects and rooms in a houseware domain, hence driving a robotic agent to choose a specific room to inspect, while solving simpler MDP problems locally. Similarly, authors of [17] propose DARLING, which uses ASP statements to bound MDP exploration in a simulated grid domain (similar to rocksample) and in a real service robotics scenario.

Logic statements can help also avoid unwanted behaviors, e.g. in safety-critic scenarios. To this aim, authors of [19, 20] refine rule templates to identify unexpected decisions and shield undesired branches of belief-action tree from exploration in POMCP, in the context of velocity regulation for a mobile robot. In [32], a goal-constrained belief space containing only safely reachable states from the initial one is defined with propositional logics, and a Satisfiability Modulo Theory (SMT) solver is used to guarantee proper execution of an houseware automation task.

Aforementioned approaches, as well as similar ones implementing more complex reasoning with, e.g., temporal logic [6, 7, 16], have one major drawback: logic statements are assumed to be provided by the user. This is unrealistic, especially in complex domains, where even expert users can hardly define accurate policy rules. Moreover, previous works use logic statements to model an additional reward in MDPs.

In contrast, we learn logic relations between domain-relevant user-defined features, which are more easily available, and actions with the paradigm of ILP under ASP semantics, and we use them to directly advise POMCP simulations. POMCP was recently used successfully in real-world applications, such as robot planning [11, 31, 33]. Learning meaningful task representations has been proposed, e.g., in [1, 26], though computational efficiency is mined by the use of an ASP meta-program. ILP has been shown to be effective in enhancing comprehensibility of black-box models [5, 25]. Recent applications in robot planning include task knowledge learning from human labelled examples with ILASP [22], and learning of temporal specifications from expert-generated MDP traces [6]. ILASP is more efficient than, e.g., [26] since it does not rely on ASP meta-program representation as of latest versions.

3 PROBLEM DEFINITION AND CASE STUDIES

The goal of this paper is to discover logical rules which underlie the decision-making process of a POMDP agent, analyzing patterns of execution. In particular, the agent performs actions depending on the current state of the environment and the goal of the assigned task. Hence, we want to find logical rules which match POMDP representation of the environment to a specific policy, with the aim of advising POMCP solver and directing the execution of the agent in future online executions. In particular, by integrating rules in POMCP exploration we aim to improve the final cumulative reward

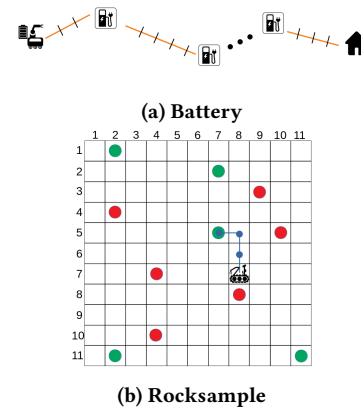


Figure 1: Example scenarios for our two case studies.

in typically challenging task instances, e.g., with long planning horizons and when the number of available actions (and corresponding needed simulations) increases. We now provide an informal description of the two domains chosen as case studies, *battery* and *rocksample*. In the following sections, we will refer to rocksample as a running example to explain the main elements of the background and methodology.

3.1 Battery

In the battery domain, an agent moves forward towards a target location, gaining a positive reward in case of success and a negative reward in case of battery depletion. The agent has limited energy autonomy, i.e., a battery with L levels of charge. Moving may reduce the level of the battery, but at known locations on the path there are recharge stations. However, recharging has a cost (i.e., a negative reward). The observable state of the system is represented by positions of the target, agent and stations, while the unobservable part of the state is the level of the battery, which can be checked with a noisy sensor. The belief in this case is a probability distribution over the possible battery levels. An explanatory scenario is depicted in Figure 1a, with segments connecting the robot, stations and the goal subdivided in discrete steps of advancement.

3.2 Rocksample

In the rocksample domain, an agent can move in cardinal directions (north, south, east and west) on a $N \times N$ grid, one cell at a time, with the goal to reach and sample a set of M rocks with known position on the grid. Rocks may be valuable or not. Sampling a valuable rock yields a positive immediate reward, while sampling a worthless rock yields negative reward. The observable state of the system is described by the current position of the agent and rocks, while the value of rocks is uncertain before sampling. The belief is in this case a probability distribution over all possible configurations of rock values. However, the agent can check the value of a rock, with an accuracy depending on the distance between the rock and the agent. Finally, the agent can obtain a positive reward exiting the grid from the right-hand side. An explanatory scenario with $N = 11$ and $M = 11$ is depicted in Figure 1b, where valuable and invaluable

rocks are marked as green and red dots, respectively, and the agent is at location (8, 7) initially.

4 BACKGROUND

We first introduce the definitions of POMDPs and POMCP, then we describe key concepts of ASP and ILP which are used in the following to describe the proposed methodology.

4.1 Partially Observable Monte Carlo Planning

A Partially Observable Markov Decision Process (POMDP) [12] is a tuple $(S, A, O, T, Z, R, \gamma)$, where S is a set of partially observable states, A is a set of actions, Z is a finite set of observations, $T: S \times A \rightarrow \Pi(S)$ is the *state-transition model*, with $\mathcal{B} = \Pi(S)$ probability distribution over states, $O: S \times A \rightarrow \Pi(Z)$ is the *observation model*, R is the *reward function* and $\gamma \in [0, 1]$ is a *discount factor*. An agent must maximize the *discounted return* $E[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$. A probability distribution over states, called *belief*, is used to represent the partial observability of the true state. To solve a POMDP it is required to find a *policy*, namely a function $\pi: \mathcal{B} \rightarrow A$ that maps beliefs \mathcal{B} into actions.

In this work, we focus on *Partially Observable Monte-Carlo Planning (POMCP)* [27] to solve POMDPs. POMCP is an *online* algorithm that solves POMDPs by using Monte Carlo techniques. The strength of POMCP is that it does not require an explicit definition of the transition model, observation model, and reward. Instead, it employs a black-box simulator of the environment. POMCP uses *Monte Carlo Tree Search (MCTS)* at each time step to explore the belief space and select the best action. *Upper Confidence Bound for Trees (UCT)* [13] is used as a search strategy to select the subtrees to explore. In particular, given history h of former belief-action pairs, UCT suggests to explore action a which maximizes the action value:

$$V_{UCT}(ha) = V(ha) + c \cdot \sqrt{\frac{\log N(h)}{N(h)}} \quad (1)$$

where $V(ha)$ is the expected return achieved by selecting action a , $N(h)$ is the number of simulations performed from history h , $N(ha)$ is the number of simulations performed while selecting action a , and c is known as the exploration constant. UCT is used to balance the exploration of new actions (i.e., action with low $N(ha)$) and the exploitation of effective actions (i.e., action with high $V(ha)$). The belief is implemented as a *particle filter*, which is a sampling over the possible states. The belief can be initialized randomly or considering prior knowledge about the environment [4]. At each time step, a particle (representing a specific state of the POMDP) is selected from the filter and used as an initial point to perform a simulation in Monte Carlo tree. The particle filter is then updated at each step after an action is performed. A simulation is a sequence of action-state pairs that collects a discounted return. Running different simulations, we can estimate and select the action which leads to the highest return. If required, new particles can be generated from the current state through a process of *particle reinvention*.

4.2 Answer Set Programming

An ASP program represents a domain of interest with a *signature* and *axioms* [2]. The signature is the alphabet of the domain,

defining its relevant attributes as variables (with discrete ranges) and predicates of variables (*atoms*). For example, in the rocksample domain variables of interest are: rock identifiers R (natural), distances between agent and rock positions D (integer), and (discretized) probabilities computed from the belief $V \in \{0, 10, \dots, 100\}$ ¹. Atoms typically represent environmental features and actions. Possible atoms in rocksample are `guess(R, V)` and `sample(R)`, denoting respectively, the probability V the rock R is valuable, and the action of sampling rock R . A variable whose value is assigned is said to be *ground* (e.g., $R=1$). An atom is ground if its variables are ground. Ground atoms become Booleans with truth value.

Axioms are logical relations between atoms. In this paper, we consider only causal rules and weak constraints. A *causal rule* $h :- b_{1..n}$ defines preconditions as the logical conjunction $\bigwedge_{i=1}^n b_i$ (*body* of the rule) for the *head* h . For instance,

$$\text{sample}(R) :- \text{guess}(R, V), V > 60. \quad (2)$$

means that rock R can be sampled if the agent believes it is valuable with probability $V > 60\%$. *Weak constraints* express preference criteria between atoms. A weak constraint has the form:

$$:\sim b_1(V_1, \dots, V_n), \dots, b_m(V_1, \dots, V_n). [w@l, V_1, \dots, V_n]$$

where w is the weight, l the (integer) priority level (used if multiple weak constraints are specified), b_i are atoms and V_j variables. The weight can either be one variable among V_j 's, or an integer. If the weight is integer, the weak constraint means that grounding of atoms b_i 's should be penalized. If the weight is a variable, say V_1 , then the constraint means that lower values for V_1 should be preferred. As an example, if (2) holds and the following weak constraint is specified:

$$:\sim \text{sample}(R), \text{guess}(R, V). [-V@1, R, V] \quad (3)$$

when ≥ 2 rocks have probability $> 60\%$ to be valuable, the one with higher V (negative weight) will be preferred.

Given an ASP task description, an ASP solver computes *answer sets*. An answer set is the minimal set of ground atoms satisfying axioms. Starting from an initial grounding, the solver deduces ground heads of rules from known ground body atoms. For the scope of this paper, answer sets contain ground features and actions. For instance, with reference to (2), if at a specific step the agent's belief distribution (formalized in ASP) is $\{\text{guess}(1, 50), \text{guess}(2, 70)\}$ (namely, the agent believes rock 1 is valuable with probability 50%, and rock 2 is valuable with probability 70%), then action `sample(2)` will be grounded, and the answer set will contain all three ground atoms, i.e., $\{\text{guess}(1, 50), \text{guess}(2, 70), \text{sample}(2)\}$.

4.3 Inductive Logic Programming

An ILP problem \mathcal{T} under ASP semantics is defined as the tuple $\mathcal{T} = \langle B, S_M, E \rangle$, where B is the *background knowledge*, i.e. a set of atoms and axioms in ASP syntax (e.g., ranges of variables); S_M is the *search space*, i.e. the set containing all possible ASP axioms that can be learned; and E is a set of *examples* (e.g., a set of ground atoms constructed from traces of execution). Our goal is to construct a set of ASP axioms belonging to S_M , called *hypothesis* H , that can explain as many as possible of the examples in E . To this end, we use

¹From now on, probability values represent percentages, e.g., $\{0\%, 10\%, \dots, 100\\}$.

state-of-the-art ILASP learner [14], where examples are *Context-Dependent Partial Interpretations* (CDPIs).

Definition 1 (Partial interpretation). Let P be an ASP program. A *partial interpretation* of P is defined as $e = \langle e^{inc}, e^{exc} \rangle$, where e^{inc} is named *included set*, i.e., a subset of ground atoms which can be part of an answer set of P ; e^{exc} is named *excluded set*, i.e., a subset of ground atoms which are not part of an answer set of P .

Definition 2 (Context-dependent partial interpretation (CDPI)). A CDPI of an ASP program P is a tuple $\langle e, C \rangle$, where e is a partial interpretation of P and C is a set of atoms called *context*.

In this paper, partial interpretations contain atoms for actions, while the context involves feature atoms. In this way, policy specifications can be learnt. In our formulation, e^{inc} includes observed actions and e^{exc} includes unobserved (hence not executed) actions. For instance, in *rocksample*, a CDPI may be

$$\langle \langle \{\text{sample}(2)\}, \{\text{sample}(1)\} \rangle, \{\text{guess}(1, 50), \text{guess}(2, 70)\} \rangle$$

where $C = \{\text{guess}(1, 50), \text{guess}(2, 70)\}$, $e^{exc} = \{\text{sample}(1)\}$, $e^{inc} = \{\text{sample}(2)\}$. The meaning of this CDPI is that, given the context defining the current belief of the agent (namely, it believes that rock 1 is valuable with probability 50% and rock 2 is valuable with probability 70%), rock 2 should be sampled and rock 1 should not.

We can now define the ILASP problem considered in this paper, hence the properties to be satisfied by H :

Definition 3 (ILASP task with CDPIs). An ILASP learning task with CDPIs is a tuple $\mathcal{T} = \langle B, S_M, E \rangle$, where E is a set of CDPIs such that:

$$\forall e = \langle \langle e^{inc}, e^{exc} \rangle, C \rangle \in E : B \cup H \cup C \models e^{inc} \wedge B \cup H \cup C \not\models e^{exc}$$

In other words, ILASP finds axioms which guarantee that actions in e^{inc} , observed in the examples, can be executed (i.e., can be grounded in an answer set), while unobserved actions in e^{exc} cannot, given the context set of environmental features. In addition, ILASP finds the *minimal* hypothesis H , i.e., axioms with the least number of atoms. This increases the comprehensibility of learned axioms and improves efficiency of ASP solving. ILASP can also learn weak constraints from *ordered CDPIs*, i.e., partial interpretations with pre-defined preference values [15].

Finally, ILASP finds the hypothesis which explains *most* of the examples, i.e., there may be CDPIs which are not covered by H . When a hypothesis is found, ILASP also returns the number of non-converged CDPIs, which quantifies the reliability of H with respect to the provided set E .

5 METHODOLOGY

We now describe our methodology for learning policy-related logic specifications from POMDP traces of execution, for a given domain of interest. In particular, we want to represent the policy map $\pi : \mathcal{B} \rightarrow \mathcal{A}$ as a set of logical formulas, through a new map $\Gamma : \mathcal{F} \rightarrow \mathcal{A}$, being $\mathcal{F} = \{F_i\}$ a set of categorical environmental features defined by a user (e.g., $\text{guess}(R, V)$ in *rocksample*), and $\mathcal{A} = \{A_i\}$ the logical formulation of A (containing, e.g., $\text{sample}(R)$ for *rocksample*). Assuming actions and features are represented as ASP atoms, three main steps are required to build map Γ and integrate logical formulas into POMCP: *ASP representation of the*

domain (Section 5.1); *definition of the ILASP problem* starting from traces of POMDP executions (Section 5.2); *integration of learned axioms in POMCP* (Section 5.3).

5.1 ASP representation of the domain

In order to represent features and actions in ASP syntax, we define a *feature map* $F_{\mathcal{F}} : \mathcal{B} \rightarrow G(\mathcal{F})$ and an *action map* $F_{\mathcal{A}} : \mathcal{A} \rightarrow G(\mathcal{A})$, being $G(\cdot)$ the *grounding function* defining all possible groundings of an atom or a set of atoms. For instance, considering $\mathcal{A} = \{\text{sample}(R)\}$ in a *rocksample* scenario with 2 rocks, $G(\mathcal{A}) = \{\text{sample}(1), \text{sample}(2)\}$. Once maps $F_{\mathcal{F}}, F_{\mathcal{A}}$ are defined by a user, it is possible to automatically translate traces of POMDP executions (i.e., sequences of belief-action pairs) to sets of ground atoms.

5.2 ILASP problem definition from POMDP traces

We can now define an ILASP task as in Section 4.3, with examples built from POMDP traces using maps $F_{\mathcal{F}}, F_{\mathcal{A}}$ and target hypothesis $H = \Gamma$. We make the following assumptions to define the ILASP learning problem:

- examples are taken from “good” traces of execution, i.e., where the policy is nearly optimal, so the agent successfully completed the task with high discounted return. In this way, learned rules are more efficient in online execution. In Section 6 we detail how POMDP traces for learning are generated in order to satisfy this assumption (e.g., tuning POMCP parameters);
- for each action, the corresponding map Γ does not depend on maps for other actions, i.e., axioms for each action are independent on axioms for other actions (e.g., rules for checking or sampling a rock are mutually independent, but are only connected to features). Hence, we can define a separate ILASP task for each action, increasing computational efficiency thanks to the shrinkage of the search space.

Each ILASP problem is defined by background knowledge B , search space S_M and examples E . In this paper, B only contains the definitions of ASP variables and ranges. Since $\Gamma : \mathcal{F} \rightarrow \mathcal{A}$, we define S_M allowing only atoms $\in \mathcal{A}$ in the head of candidate axioms and atoms $\in \mathcal{F}$ in the body. Examples are extracted from traces of POMDP executions using maps $F_{\mathcal{F}}$ and $F_{\mathcal{A}}$. Specifically, for each belief-action $\langle b, a \rangle \in \mathcal{B} \times \mathcal{A}$, we obtain a pair of ground feature set and action $\langle \text{belief}, \text{action} \rangle$, where $\text{belief} \subseteq G(\mathcal{F})$ and $\text{action} \subseteq G(A_i)$, A_i being the ASP atom representing a . For instance, in a *rocksample* instance with 2 rocks, assume sampling of the second rock is performed, and the probabilities that rocks 1 and 2 are valuable are 70% and 80%, respectively. Then, $A_i = \text{sample}(R)$, $\text{action} = \text{sample}(2)$ and $\text{belief} = \{\text{guess}(1, 70), \text{guess}(2, 80)\}$. We then generate a CDPI in the form

$$\langle \langle \{\text{action}\}, G(A_i) \setminus \{\text{action}\} \rangle, \text{belief} \rangle \quad (4)$$

where $G(A_i) \setminus \{\text{action}\} = \{\text{sample}(1)\}$ in the previous example. Moreover, at each step where an action a is not executed, we define an additional CDPI

$$\langle \langle \emptyset, G(A_i) \rangle, \text{belief} \rangle \quad (5)$$

For instance, in the previous example, if a is the checking action, $G(A_i) = \{\text{check}(1), \text{check}(2)\}$, and similarly for other actions. In this way, learnt axioms will provide more meaningful and useful policy specifications, since they will be induced also from *counterexamples* (i.e., examples where actions are not executed).

In order to learn weak constraints (i.e., preferences on specific actions with respect to others), for each example as (4) we generate CDPIs in the form:

$$\langle\langle G(A_i) \setminus \{\text{action}\}, \emptyset \rangle, \text{belief} \rangle \quad (6)$$

In other words, all non-executed actions are in e^{inc} , while $e^{exc} = \emptyset$. We then specify a preference of (4) over (6). This is needed, since execution traces only contain a single action at each step, but possibly others could be executed. For instance, consider the former rocksample instance and assume (2) holds. Sampling of both rocks is possible, since they are valuable with probability $\geq 60\%$, but only the second rock was actually sampled. We then define the following examples, corresponding respectively to (4)-(6):

$$\langle\langle \{\text{sample}(2)\}, \{\text{sample}(1)\} \rangle, \{\text{guess}(1, 70), \text{guess}(2, 80)\} \rangle \quad (7a)$$

$$\langle\langle \{\text{sample}(1)\}, \emptyset \rangle, \{\text{guess}(1, 70), \text{guess}(2, 80)\} \rangle \quad (7b)$$

Then, we specify that (7a) must be preferred to (7b).

For each action, we consider hypothesis which covers most of the examples and use it to bias POMCP exploration, as explained in next subsection.

5.3 Integrating axioms in POMCP

We integrate the computed axioms inside POMCP to help the algorithm selecting the best action given the current state and belief. Specifically, when a new node is reached during the Monte Carlo Tree Search (MCTS), we instantiate the variables of the axioms with values computed from the current belief. POMCP represents the belief of the root node using a particle filter (i.e., the algorithm keeps a finite collection of possible states for each belief). Thus, we compute the values required to ground ASP features by analyzing all the particles. For example, to compute $\text{guess}(R, V)$ for rock 1 we compute the percentage of particles in which rock 1 is valuable. When we reach a new non-root node, no particle filter is available. However, we still ground features that depend only on the observable part of the state (e.g., the current position in rocksample). These features need to be re-grounded only for previously visited non-root nodes that become a root (i.e., we move the simulation one step forward by selecting an action and receiving an observation, and the new root is a node that we visited previously during a simulation). In this case, the features are re-instantiated using the new full belief. In all the other cases, no recomputation is needed since the observable part of a node never changes. When all features of a node are instantiated, we call the ASP solver to ground new actions.

To combine logical rules with POMCP, we introduce a prior in UCT (Equation 1). Specifically, for each action a grounded by the ASP solver, we increase the values $V(ha)$ to the same value of c , and $N(ha)$ to a fixed value representing high reward (this is domain-dependent, we empirically set it to 10 in our domains). This is the same as implying that we have already performed $N(ha)$ simulations in the current belief, and they achieved a good return value (i.e., an high value of $V(ha)$). It is important to notice that

this prior does not impact the optimality of the policy. In fact, the number $N(ha)$ is finite, thus, as in standard POMCP, performing an infinite number of simulations we would converge toward the optimal policy as in standard POMCP. However, this prior considerably decreases the need for exploration during MCTS, since the most probable best actions are suggested by axioms before MCTS simulation starts. This is a crucial benefit when the tree grows in size (hence, many simulation are needed), as shown in Section 6.

6 EMPIRICAL EVALUATION

We now describe the empirical evaluation for our two case studies, rocksample and battery.

6.1 Empirical methodology

Learning executions are generated running POMCP in 1000 random scenarios for each domain, specifically considering: 12X12 grid with 4 rocks for rocksample, randomizing the value and position of rocks and initial position of the agent; 35-step path for battery, randomizing locations of stations with a limited range of mutual distances ($\in [1, 4]$). We set 2^{15} particles² and simulations in POMCP algorithm, in order to generate “good” executions as explained in Section 5.2. ILASP examples are extracted only from execution traces with discounted return greater or equal than the average over all traces. This criterion is domain-independent to select only meaningful examples to learn from.

We evaluate performance of POMCP with and without the bias of ASP axioms in two scenarios: i) using fewer particles with respect to the training set, while keeping problem dimension fixed (50 random settings for each number of particles are generated); ii) increasing the problem dimension with fixed number of particles to 2^{15} . Specifically, in the second case we increase the grid size N in rocksample and path length in battery, running 50 random simulations for each parameter value. In this way, we show the scalability of rules to different instances of the domains, and their utility in practical applications, where the planning horizon is critical for POMCP exploration. ASP solving is performed with state-of-the-art Clingo [9]. All experiments are run on a notebook with Intel Core i7-6700HQ and 16GB of RAM.

6.2 Battery

Results for battery domain for each step of our methodology (Section 5) follow.

6.2.1 ASP representation. We define the following features for the battery domain:

- $\text{guess}(L, V)$, representing the probability $V \in \{0, 10, \dots, 100\}\%$ that the level of the battery is at least $L \in \{0, 1, \dots, 10\}$.
- $\text{dist_next}(D)$, representing the distance D (integer) from the next station, if any, or the goal.
- $X \geq \bar{x}$ and $X \leq \bar{x}$, where X is either L, D, V and \bar{x} is a possible value for X .
- at_station , meaning that the agent is on a station location.

Atoms for actions are recharge, check, advance.

²As of standard practice, the number of particles is the same as the number of simulations.

6.2.2 *ILASP results.* ILASP learns the following causal rules from a set of $\approx 29600^3$ examples (i.e., CDPIs as in Section 4.3):

```
check :- V ≥ 20; L ≤ 4; guess(L, V).
advance :- dist_next(D); D ≤ 4.
recharge :- L ≤ 7; V ≥ 30; D ≥ 2; dist_next(D);
           guess(L, V); at_station.
```

Checking is convenient if the level of battery is considerably low (≤ 4) with non-negligible probability ($\geq 20\%$). Recharging should be performed if the level of battery is ≤ 7 with good probability ($\geq 30\%$), and the next station is not sufficiently close ($D \geq 2$). Finally, advancing is always suggested, since the next station (or the goal) is always closer than 4 steps in our problem definition. The learning time is ≈ 45 s for the longest recharge rule (requiring more time), with S_M including ≈ 4700 axioms⁴.

6.2.3 *POMCP integration results.* In Figure 2a, we show the discounted return achieved by POMCP (with and without rule bias) when the number of particles is varied between 2^8 and 2^{15} . Introducing learned ASP axioms does not significantly affect the performance of the solver. This is probably due to the limited number of actions involved in the task (only 3, i.e., advance, check and recharge), thus few branches must be explored at each simulation step.

Figure 2b shows the discounted return achieved by POMCP when the length of the path increases, while the number of particles is constantly 2^{15} as in the training setup. In this case, the solver performs significantly better when learned rules are employed (the discounted return increases up to 20% when the path length is 75, i.e., almost double than the training scenario). In fact, as the path length increases, the risk of battery discharge is higher, since failing to recharge is more probable. Moreover, the planning horizon, hence simulation depth, also increases, so guidance of ASP axioms is crucial.

6.3 Rocksample

Results for rocksample domain for each step of our methodology (Section 5) follow.

6.3.1 *ASP representation.* We define the following set of features:

- $\text{guess}(R, V)$, representing the probability V that rock R is “good”. Since ASP can only deal with integer variables, V is discretized in the range $\{0, 10, \dots, 100\}$ with the following meaning: $\text{guess}(1, 10)$ means that rock 1 is valuable with probability $P(1) \in [10, 20[$ %. We choose this discretization step for probabilities empirically; finer discretization is possible, but may affect ASP computational time in large domains.
- $\text{dist}(R, D)$, representing the 1-norm distance D between the agent and rock R .
- $\text{min_dist}(R)$, meaning that rock R is the closest to the agent.
- $\text{delta_x}(R, D)$ (respectively, $\text{delta_y}(R, D)$), meaning x - (respectively, y -)coordinate of rock R with respect to the agent is D .

³The number of examples depends on the number of steps in executions.

⁴The size of S_M depends on the number of features, variables and their possible values.

- $X \geq \bar{x}$ and $X \leq \bar{x}$, where X is either V, D and \bar{x} is a possible value for X .
- $\text{sampled}(R)$, meaning that a rock R has been sampled.
- $\text{num_sampled}(N)$, meaning that percentage N of rocks has been sampled ($N \in \{0, 25, \dots, 100\}$).
- $\text{target}(R)$, meaning that rock R is the next to be sampled.

In particular, min_dist and best_guess are qualitative features which are added to possibly simplify ASP rule expressions and improve their comprehensibility.

Action atoms are $\text{east}, \text{west}, \text{north}, \text{south}, \text{exit}^5, \text{check}(R), \text{sample}(R)$, where argument R represents the target rock to be checked or sampled.

6.3.2 *ILASP results.* ILASP learns the following axioms from a set of ≈ 8500 examples:

$$\text{east} :- \text{target}(R), \text{delta_x}(R, D), D \geq 1. \quad (8a)$$

$$\text{west} :- \text{target}(R), \text{delta_x}(R, D), D \leq -1. \quad (8b)$$

$$\text{north} :- \text{target}(R), \text{delta_y}(R, D), D \geq 1. \quad (8c)$$

$$\text{south} :- \text{target}(R), \text{delta_y}(R, D), D \leq -1. \quad (8d)$$

$$\text{target}(R) :- \text{dist}(R, D), \text{not sampled}(R), D \leq 1. \quad (8e)$$

$$\text{target}(R) :- \text{guess}(R, V), \text{not sampled}(R), \\ 70 \leq V \leq 80. \quad (8f)$$

$$\text{check}(R) :- \text{target}(R), \text{not sampled}(R), \\ \text{guess}(R, V), V \leq 50. \quad (8g)$$

$$\text{check}(R) :- \text{guess}(R, V), \text{not sampled}(R), \\ \text{dist}(R, D), D \leq 0, V \leq 80. \quad (8h)$$

$$\text{sample}(R) :- \text{target}(R), \text{dist}(R, D), D \leq 0, \\ \text{not sampled}(R), \text{guess}(R, V), \\ V \geq 90. \quad (8i)$$

$$\text{exit} :- \text{guess}(R, V), V \leq 40, \text{not sampled}(R), \\ \text{num_sampled}(N), N \geq 25. \quad (8j)$$

$$\text{exit} :- \text{dist}(R, D), 5 \leq D \leq 8, \text{not sampled}(R), \\ \text{num_sampled}(N), N \geq 25. \quad (8k)$$

Axioms for moving actions are quite simple. Motion depends on the target rock and relative position with respect to it (delta_x and delta_y atoms). Two axioms are found for target : an unsampled rock is chosen as a target if either its distance from the agent is low (≤ 1) or it is valuable with high probability ($\text{guess } V$ between 70% and 80%). Sampling is then performed on the target rock if the agent is at its location and the probability to be valuable is $\geq 90\%$. Two axioms were discovered also for check , which occurs either when the target rock has low ($\leq 50\%$) probability to be valuable, or to verify the value of a rock when on it. Finally, the agent decides to exit the grid when at least 25% of rocks have been sampled and either an unsampled rock has low probability to be valuable ($\leq 40\%$) or it is far from the agent (distance between 5 and 8).

⁵ exit is considered as a separate action for generality. Since the agent exits the grid from the right-hand side in our domain, it is manually mapped to east action for the scope of this paper, in POMDP representation.

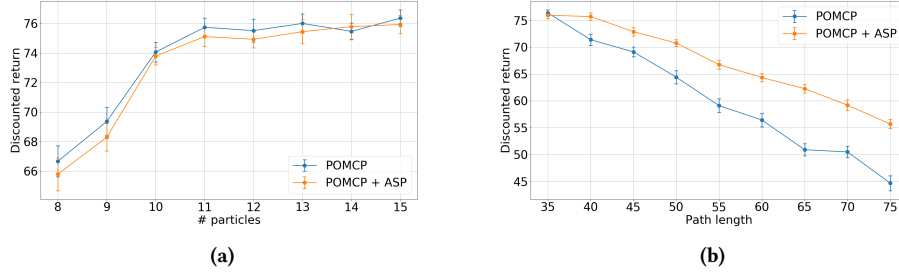


Figure 2: Battery (mean \pm std dev). a. Tests with reduced number of POMCP particles; b. with increasing path length.

ILASP is also able to learn the following set of weak constraints for the target atom:

```

: ~ target(R), dist(R,D).[D@1, R, D]
: ~ target(R), min_dist(R), guess(R,V).
[-V@2,R,V]

```

In other words, when multiple rocks can be targets according to causal rules, the agents ranks them considering primarily the closeness, secondly the probability to be valuable. The learning time is ≈ 420 s for the longest target rules (with weak constraints), with S_M including ≈ 29500 axioms.

6.3.3 POMCP integration results. Figure 3a shows the performance of POMCP (with and without rule bias) when the number of particles varies between 2^8 and 2^{15} with grid size and number of rocks as in the training setup ($N = 12$ and $M = 4$). The discounted returns achieved with and without the rules are very similar to each other when a low number of particles (i.e., $\leq 2^{11}$) is used. When instead the number of particles increases (i.e., $> 2^{11}$), the learned rules improve the performance (up to 15 – 20% with 2^{15} particles)⁶. In fact, the rules describe more accurately the policy map when a larger amount of particles is used, because the precision of the belief distribution is higher, hence, upper and lower bounds about the distance and probability of rocks to be valuable are more precise (e.g., $D \leq 0$ and $V \geq 90$ in (8i)). Since the number of actions in rocksample is higher than the battery domain, learned axioms significantly reduce the number of branches to be expanded in the Monte Carlo tree in POMCP exploration. This is even more evident in Figure 3b, where the grid size is increased from 12 to 24, while the number of rocks and simulations are kept equal to those in the training set, i.e., $M = 4$ and 2^{15} respectively. In this case, the planning horizon grows, and rules become crucial to achieve better performance. In fact, POMCP with rules reaches almost double discounted return than standard POMCP, on average, with a 24×24 grid.

6.4 Further considerations

We now analyze some crucial aspects of our pipeline, specifically the computational impact of answer set solving in POMCP (Section 6.4.1) and how the quantity and quality of examples in the training set of executions affect learning and planning outcomes (Sections 6.4.2–6.4.3). For brevity, we consider only the rocksample scenario,

⁶The dip at 2^{14} particles is due to the lower number of valuable rocks, on average, in the random scenarios.

which turned out to be more challenging than battery, since it has more actions.

6.4.1 Computational impact of ASP in POMCP. As explained in Section 5.3, ASP solving occurs at every node in MCTS. In the rock-sample scenario with 2^{15} particles and 4 rocks on a 12×12 grid, Clingo is invoked on average 125000 times per trial, over 50 task instances. Each ASP solving call takes approximately 0.3 ms. Compared to pure POMCP solver in the same scenario, this results in higher computational time per step, i.e., executed action (5.53s vs. 1.88s on average) and over all task instance (27s vs. 18s on average). Using Clingo has the advantage to keep the implementation modular, so it is easy to modify task axioms in a separate ASP file to be evaluated at runtime. However, it is more efficient to implement ASP axioms directly as conditional statements in POMCP code (an example is provided in the linked repository), resulting in similar computational time as pure POMCP (20s vs. 18s on average per task instance, in previous conditions). Moreover, Figure 3a shows that axioms allow to achieve higher discounted return for $\geq 2^{12}$ particles than pure POMCP with $N = 2^{15}$. POMCP+Clingo completes a task instance in ≈ 2.5 s on average with 2^{12} particles, almost one order of magnitude less than pure POMCP with 2^{15} particles. Thus, our pipeline can obtain better performance with lower computational cost.

6.4.2 Robustness of learned axioms. We now evaluate how the number of examples in the training set for ILASP affects learned axioms. Figures 4a–4b show the results about rule learning depending on the number of examples (i.e., CDPIs) used. In the x -axis we have the percentage of examples with respect to the original training set used in previous experiments (i.e., ≈ 8500 in rocksample). In the y -axis of Figure 4a we show the *percentage of non-covered examples*, while in the y -axis of Figure 4b we show the *distance between learned axioms* from a given amount of examples and from the full training set (normalized by the number of atoms in the latter). Given 2 rules R_1, R_2 , each one made of a set of atoms $\{a_i\}, i \in \{1, 2\}$, we define their distance as $R_1 - R_2 = |\{a_1\} \cup \{a_2\}| - |\{a_1\} \cap \{a_2\}|$, with $|\cdot|$ the cardinality of a set. For instance, $\text{sample}(R) :- \text{dist}(R,V), V \leq 2$ has a distance 5 from (8i), due to the missing $\text{not sampled}(R), \text{guess}(R,V), V \geq 90$ and $\text{target}(R)$, and the different upper bound on distance D . We observe that using $\geq 80\%$ of the full training dataset, the percentage of non-covered examples stabilizes (Figure 4a) and the distance becomes null for all actions (Figure 4b), thus learned rules




Figure 3: Rocks sample (mean ± std dev). a. Tests with reduced number of POMCP particles; b. with enlarged grid size N .

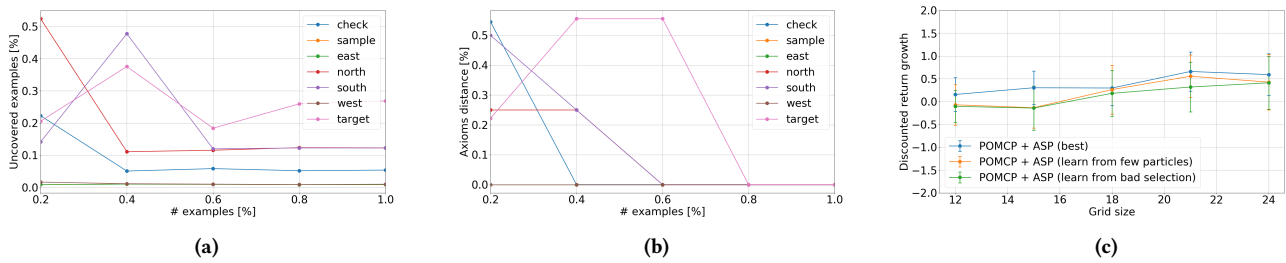


Figure 4: a. Non-covered examples (%) for different number (#) of examples used for learning; b. Relative distance between rules learned with different # of examples, w.r.t. rules learned from the full training set; c. Growth rate of discounted return (mean ± std deviation) for rocks sample with larger grid size, when axioms are learned from best examples (2^{15} particles and discounted return > average) or sub-optimal examples (either discounted return < average or generated from POMCP with 2^{11} particles).

successfully converge to stable hypotheses. Overall, rules learned from the full dataset cover more > 73% of examples.

6.4.3 *Influence of quality of examples on POMCP performance.* We now evaluate the influence of the quality of examples (i.e., of the policy generating them) on the discounted return achieved by POMCP with ASP axioms. We want to show what happens when an optimal policy cannot be computed by POMCP for training set generation due to the complexity of the task, so only a sub-optimal one is available. For testing, we first learn rules for rocksample from POMCP optimal executions generated with 2^{15} particles, but selecting only traces with discounted return *smaller* than the average on all traces. This emulates the case in which a suitable criterion for example selection cannot be easily defined, e.g., in complex domains. Then, in another test we learn rules from traces generated with 2^{11} particles, i.e., where axioms improve performance of pure POMCP, and select traces with discounted return > average. This is also useful for fast generation of examples in complex domains. Figure 4c shows the results for the two tests (considering 50 random settings for each value of grid size), in terms of the rate of *growth of discounted return*, i.e., the relative difference between discounted return achieved with and without rules, normalized by the discounted return achieved without rules. In our charts, POMCP is run with 2^{15} particles when combined with rules for validation. In general, rules learned from sub-optimal examples do not degrade performance with respect to pure POMCP (only a slight average decrease < 14% is observed for grid size $N < 18$), thanks to the soft guidance approach. Interestingly, sub-optimal ASP rules learned in both tests still are beneficial

to POMCP as the planning horizon increases ($N \geq 18$). In particular, at $N = 24$ discounted return increases of approximately 50%.

7 CONCLUSION

We presented a novel methodology for soft POMCP policy guidance, based on ILP to learn ASP rules describing policy specifications directly from POMDP execution traces. Our rules do not need to be hand-crafted by experts, but we only require definition of relevant high-level domain-specific features, in order to describe belief information in ASP formalism and find matches with POMDP actions. We showed that learned ASP axioms significantly support POMCP exploration in the benchmark domains of rocksample and battery, advising good branches to be explored in POMCP simulation without significantly affecting the computational burden. In particular, rules are more beneficial when the number of possible actions increases (rocksample) and the planning horizon is extended. Furthermore, POMCP property of asymptotic optimality is preserved when axioms are learned from traces generated with low-quality policies, and axioms converge to a fixed point as the number of examples increases.

In the future, we plan to investigate i) how to integrate more advanced (e.g., temporal) logic specifications to address more complex tasks; ii) how to account for non-covered examples in the training set and possibly refine axioms during execution; iii) the influence of bad definition of environmental features by users.

ACKNOWLEDGMENTS

This project has received funding from the Italian Ministry for University and Research, under the PON “Ricerca e Innovazione” 2014-2020” (grant agreement No. 40-G-14702-1).

REFERENCES

- [1] Blai Bonet and Hector Geffner. 2020. Learning First-Order Symbolic Representations for Planning from the Structure of the State Space. In *ECAI 2020*. IOS Press, 2322–2329.
- [2] Francesco Calimeri, Wolfgang Faber, Martin Gebser, Giovambattista Ianni, Roland Kaminski, Thomas Krennwallner, Nicola Leone, Marco Maratea, Francesco Ricca, and Torsten Schaub. 2020. ASP-Core-2 input language format. *Theory and Practice of Logic Programming* 20, 2 (2020), 294–309.
- [3] Anthony R Cassandra, Michael L Littman, and Nevin Lianwen Zhang. 2013. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. *arXiv preprint arXiv:1302.1525* (2013).
- [4] Alberto Castellini, Georgios Chalkiadakis, and Alessandro Farinelli. 2019. Influence of State-Variable Constraints on Partially Observable Monte Carlo Planning. In *IJCAI 2019, Macao, China, August 10-16, 2019*. ijcai.org, 5540–5546.
- [5] Fabio AD’Asaro, Matteo Spezialetti, Luca Raggioli, and Silvia Rossi. 2020. Towards an inductive logic programming approach for explaining black-box preference learning systems. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, Vol. 17. 855–859.
- [6] Giuseppe De Giacomo, Marco Favorito, Luca Iocchi, and Fabio Patrizi. 2020. Imitation learning over heterogeneous agents with restraining bolts. In *Proceedings of the international conference on automated planning and scheduling*, Vol. 30. 517–521.
- [7] Giuseppe De Giacomo, Luca Iocchi, Marco Favorito, and Fabio Patrizi. 2019. Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications. In *Proceedings of the international conference on automated planning and scheduling*, Vol. 29. 128–136.
- [8] Esra Erdem and Volkan Patoglu. 2018. Applications of ASP in robotics. *KI-Künstliche Intelligenz* 32, 2 (2018), 143–149.
- [9] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, Max Ostrowski, Torsten Schaub, and Philipp Wanko. 2016. Theory solving made easy with clingo 5. In *Technical Communications of the 32nd International Conference on Logic Programming (ICLP 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [10] Michele Ginesi, Daniele Meli, Andrea Roberti, Nicola Sansonetto, and Paolo Fiorini. 2020. Autonomous task planning and situation awareness in robotic surgery. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3144–3150.
- [11] Francesco Giuliani, Alberto Castellini, Riccardo Berra, Alessio Del Bue, Alessandro Farinelli, Marco Cristani, Francesco Setti, and Yiming Wang. 2021. POMP++: Pomcp-Based Active Visual Search in Unknown Indoor Environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Prague, Czech Republic). IEEE Press, 1523–1530.
- [12] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and Acting in Partially Observable Stochastic Domains. *Artif. Intell.* 101, 1–2 (May 1998), 99–134.
- [13] Levente Kocsis and Csaba Szepesvári. 2006. Bandit Based Monte-Carlo Planning. In *Proc. ECML’06*. Springer-Verlag, Berlin, Heidelberg, 282–293.
- [14] Mark Law, Alessandra Russo, and Krysia Broda. 2015. The ILASP system for learning Answer Set Programs. www.ilasp.com.
- [15] Mark Law, Alessandra Russo, and Krysia Broda. 2016. Iterative learning of answer set programs from context dependent examples. *Theory and Practice of Logic Programming* 16, 5-6 (2016), 834–848.
- [16] Matteo Leonetti, Luca Iocchi, and Fabio Patrizi. 2012. Automatic generation and learning of finite-state controllers. In *International Conference on Artificial Intelligence: Methodology, Systems, and Applications*. Springer, 135–144.
- [17] Matteo Leonetti, Luca Iocchi, and Peter Stone. 2016. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artificial Intelligence* 241 (2016), 103–130.
- [18] Vladimir Lifschitz. 1999. Answer set planning. In *International Conference on Logic Programming and Nonmonotonic Reasoning*. Springer, 373–374.
- [19] Giulio Mazzi, Alberto Castellini, and Alessandro Farinelli. 2021. Identification of Unexpected Decisions in Partially Observable Monte-Carlo Planning: A Rule-Based Approach. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 889–897.
- [20] Giulio Mazzi, Alberto Castellini, and Alessandro Farinelli. 2021. Rule-based Shielding for Partially Observable Monte-Carlo Planning. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, Vol. 31. 243–251.
- [21] Daniele Meli, Hirenkumar Nakawala, and Paolo Fiorini. 2023. Logic programming for deliberative robotic task planning. *Artificial Intelligence Review* (2023), 1–39.
- [22] Daniele Meli, Mohan Sridharan, and Paolo Fiorini. 2021. Inductive learning of answer set programs for autonomous surgical task planning. *Machine Learning* 110, 7 (2021), 1739–1763.
- [23] Stephen Muggleton. 1991. Inductive logic programming. *New generation computing* 8, 4 (1991), 295–318.
- [24] Christos H Papadimitriou and John N Tsitsiklis. 1987. The complexity of Markov decision processes. *Mathematics of operations research* 12, 3 (1987), 441–450.
- [25] Johannes Rabold, Michael Siebers, and Ute Schmid. 2018. Explaining black-box classifiers with ILP—empowering LIME with Aleph to approximate non-linear decisions with relational rules. In *International Conference on Inductive Logic Programming*. Springer, 105–117.
- [26] Ivan D Rodriguez, Blai Bonet, Javier Romero, and Hector Geffner. 2021. Learning First-Order Representations for Planning from Black Box States: New Results. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, Vol. 18. 539–548.
- [27] David Silver and Joel Veness. 2010. Monte-Carlo planning in large POMDPs. *Advances in neural information processing systems* 23 (2010).
- [28] Adhiraj Somani, Nan Ye, David Hsu, and Wee Sun Lee. 2013. DESPOT: Online POMDP planning with regularization. *Advances in neural information processing systems* 26 (2013).
- [29] Mohan Sridharan, Michael Gelfond, Shiqi Zhang, and Jeremy Wyatt. 2019. REBA: A refinement-based architecture for knowledge representation and reasoning in robotics. *Journal of Artificial Intelligence Research* 65 (2019), 87–180.
- [30] Eleonora Tagliabue, Daniele Meli, Diego Dall’Alba, and Paolo Fiorini. 2022. De-liberation in autonomous robotic surgery: a framework for handling anatomical uncertainty. In *2022 IEEE International Conference on Robotics and Automation (ICRA 2022)*.
- [31] Yiming Wang, Francesco Giuliani, Riccardo Berra, Alberto Castellini, Alessio Del Bue, Alessandro Farinelli, Marco Cristani, and Francesco Setti. 2020. POMP: Pomcp-based Online Motion Planning for active visual search in indoor environments. In *Proc. of British Machine Vision Conference (BMVC)*.
- [32] Yue Wang, Abdullah Al Redwan Newaz, Juan David Hernández, Swarat Chaudhuri, and Lydia E Kavrakı. 2021. Online partial conditional plan synthesis for POMDPs with safe-reachability objectives: Methods and experiments. *IEEE Transactions on Automation Science and Engineering* 18, 3 (2021), 932–945.
- [33] Maddalena Zuccotto, Marco Piccinelli, Alberto Castellini, Enrico Marchesini, and Alessandro Farinelli. 2022. Learning State-Variable Relationships in POMCP: A Framework for Mobile Robots. *Frontiers in Robotics and AI* 9 (2022).