

Answer Set Programming for Non-Stationary Markov Decision Processes

Leonardo A. Ferreira · Reinaldo A. C. Bianchi · Paulo E. Santos · Ramon Lopez de Mantaras

Received: date / Accepted: date

Abstract Non-stationary domains, where unforeseen changes happen, present a challenge for agents to find an optimal policy for a sequential decision making problem. This work investigates a solution to this problem that combines Markov Decision Processes (MDP) and Reinforcement Learning (RL) with Answer Set Programming (ASP) in a method we call ASP(RL). In this method, Answer Set Programming is used to find the possible trajectories of an MDP, from where Reinforcement Learning is applied to learn the optimal policy of the problem. Results show that ASP(RL) is capable of efficiently finding the optimal solution of an MDP representing non-stationary domains.

Keywords Non-eterminism · Markov Decision Processes · Answer Set Programming · Action Languages

1 Introduction

John McCarthy defined Elaboration Tolerance as “the ability to accept changes to a person’s or a computer program’s representation of facts about a subject without having to start all over” [17]. An example of a real world problem that

Leonardo Anjoletto Ferreira
Universidade Metodista de São Paulo, Rua Alfeu Tavares, 149, São Bernardo do Campo, São Paulo, Brazil. E-mail: leonardo.ferreira@metodista.br

Reinaldo Augusto da Costa Bianchi
Centro Universitário da FEI, Av. Humberto de Alencar Castelo Branco, 3972, São Bernardo do Campo, São Paulo, Brazil. E-mail: rbianchi@fei.edu.br

Paulo Eduardo Santos
Centro Universitário da FEI, Av. Humberto de Alencar Castelo Branco, 3972, São Bernardo do Campo, São Paulo, Brazil. E-mail: psantos@fei.edu.br

Ramon Lopez de Mantaras
Institut d’Investigació en Intel·ligència Artificial, 08193 Bellaterra, Catalonia, Spain E-mail: mantaras@iiia.csic.es

requires solutions that are tolerant to elaborations is the dynamics of urban mobility, where streets and roads are constantly reconstructed or modified. Some of these changes are planned and, thus, can be previously informed to the inhabitants of the city. However, unplanned changes due to natural phenomena (rain or snowing, for example), or due to human actions (e.g. road accidents), may occur that cause road blocks which could prevent the traffic through certain routes of the city. In such cases, it is not possible to know the changes until they are observed by the agents. However, an agent immersed in this domain must be capable of finding the best sequence of actions, considering the new situations, but without losing all the information previously acquired.

One formalism that can be used to model the kind of situations described above is a non-stationary Markov Decision Process (MDP), where the set of states represented by observations of the environment (facts) can suffer changes over time such that states can be added to, or removed from, the decision process. As these changes may not be known *a priori*, the environment cannot be modelled as a stationary MDP due to the Curse of Dimensionality [9], which describes the growth in the set of states when considering the number of variables involved in the description of a state.

This work is directed towards problem solving in non-stationary domains in which, not only the transition and the reward functions change, but also the states and actions may change during the agent's interaction with the environment. The ASP(RL) proposed here is able to change an MDP's description during learning and to reuse the learnt data in the new domain that it is interacting with. A consequence of using ASP(RL) is the speed up in the searching for an MDP solution as a consequence of the reduction that may occur in the search space.

In order to model an agent capable of interacting efficiently with non-stationary domains, we propose a method called ASP(RL) that combines Markov Decision Process, Reinforcement Learning (RL) (Section 2.1) with Answer Set Programming (Section 2.2). The proposed combination (Section 3) allows an agent to learn incrementally in an environment that suffers changes. The method was analysed in a non-stationary grid world (Section 4) and experimentally evaluated and compared to two Reinforcement Learning algorithms (Section 5).

2 Background

This section introduces Markov Decision Processes (MDP), Reinforcement Learning (RL) and Answer Set Programming (ASP), which constitute the foundations of this work.

2.1 MDP and Reinforcement Learning

In a Sequential Decision Making Problem, an agent must select a series of actions in order to find a solution to a given problem. A feasible solution,

known as policy (π), is a sequence of non-deterministic actions that leads the agent from an initial state to a goal state [7,9]. A problem such as this may have more than one feasible solution, thus it is possible to use the Bellman's Principle of Optimality [7,9] as a criterion to define which of the feasible policies can be considered as the optimal policy (π^*). Bellman's Principle of Optimality states that *"an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision"* [9]. By this definition, an optimal policy is the one that maximises (or minimises) a desired reward/cost function.

Markov Decision Process (MDP) [8] can be used to formalise Sequential Decision Making Problems. An MDP is defined as a tuple $M = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ where:

- \mathcal{S} is the set of states at any time step;
- \mathcal{A} is the set of allowed actions in the states $s \in \mathcal{S}$;
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ is the transition function that gives the probability of reaching the future state $s' \in \mathcal{S}$ by performing action $a \in \mathcal{A}$ in the current state $s \in \mathcal{S}$;
- $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$ is the reward function that returns a real value for reaching a state $s' \in \mathcal{S}$ after performing an action $a \in \mathcal{A}$ in a state $s \in \mathcal{S}$.

To find the optimal solution of an MDP is to find, for each state, which is the action that maximises the reward function. One of the methods that can be used to approximate such optimal solution is Reinforcement Learning (RL). With RL, at each time step, a learning agent at a state $s \in \mathcal{S}$ chooses an action $a \in \mathcal{A}$ to be performed in the environment. After the action a is performed, the agent receives its new state $s' \in \mathcal{S}$ and a reward $r(s, a, s')$. This reward is used to update a value function $V(s)$ (or an action-value function $Q(s, a)$, depending on the method used) and the interaction continues from the new state. Given enough time, the agent is capable to approximate the (action-) value function, maximising the reward function and finding the optimal policy. One important aspect of RL methods is that the transition and reward function are not necessarily known beforehand by the agent, but are present in the environment.

Two well-known methods of RL are SARSA [23] and Q-Learning [24,23]. Both are based on the concept of updating an (action-) value function considering the observations received from the environment. The main difference between them is how this update is accomplished. SARSA is an *on-policy* method, which means that updates in the $Q(s, a)$ function use the actions executed in the policy that is being followed, while Q-Learning is an *off-policy* method that uses the maximum value of the next state to update the current state-action pairs.

Although Reinforcement Learning allows for learning the optimal solution of a sequential decision-making problem with non-stationary transition and reward functions (functions that may change over time) and without the knowledge of the reward function, it still needs stationary sets (which do not

change during the interaction) of states and actions in order to proceed with the learning process. In order to account for changes in the set of states, we propose the use of Answer Set Programming.

2.2 Answer Set Programming

Answer Set Programming (ASP) is a declarative non-monotonic logic programming language that has been used with great success to describe and provide solutions for NP-complete problems, such as planning and scheduling [15,25]. Furthermore, ASP can be used for problems with large search space, such as the Reaction Control System of a Space Shuttle [3,5,4,18].

An ASP program is a set of rules, each rule is composed of an atom A and of literals L_m , which are atoms or negated atoms. An ASP rule can be represented as: $A \leftarrow L_1, \dots, L_n$; where A is called the head of the rule and the conjunction of literals L_1, \dots, L_n is its body. A rule is said to be positive when there is no negated atom in its body; when $n = 0$ the atom A is said to be a fact.

Let Π be an ASP program, an answer set of Π is an interpretation that makes all the rules of this program true. This interpretation is a minimal model of the program. One important aspect of ASP is its non-monotonic semantics (based on the Stable Model Semantics [13]), which respects the rationality principle that states that “*one shall not believe anything one is not forced to believe*” [13]. Along with *true* and *false*, ASP also has a third truth value for *unknown*.

There are two types of negation in ASP: strong (or “classical”) and weak, which in ASP represents *negation as failure* [16].

Given an ASP program Π and a set M of atoms of Π , a reduct program Π_M is obtained from Π by [13]:

- Deleting each rule with a negative literal in its body in the form $\neg B, B \in M$;
- Deleting every negative literal in the body of remaining rules.

Thus, the reduct program Π_M is negation-free and has a unique minimal Herbrand model. If M coincides with this model for Π_M , then M is a stable model of Π . Furthermore, by using an operator O_Π defined as “for any set of atoms M of Π , $O_\Pi(M)$ is the minimal Herbrand model of Π_M ”, then a stable model can also be described as the fixed points of O_Π . From this definition, a minimal model that accepts classical negation is called an answer set instead of a stable model.

Although ASP does not provide syntax to describe non-deterministic events, it is possible to use choice rules in order to verify each possible outcome of a choice. Considering for example that an agent is at a state **s0** and chooses to perform action **a** with the possible outcomes being the future states **s1**, **s2** and **s3**, this transition can be encoded using “**1 { s1, s2, s3 } 1 :- s0, a.**” in an ASP program. Thus, when **s0** and **a** are true in Π (the agent has

performed the action \mathbf{a} in the state \mathbf{s}_0), only one of the future states \mathbf{s}_1 , \mathbf{s}_2 or \mathbf{s}_3 is true (reached by the agent).

Since ASP can be used as a tool for providing reasoning and knowledge revision on a set of states and Reinforcement Learning allows for learning the solution of an MDP without the need of an explicit reward function, an opportunity arises to combine both methods in order to efficiently find the optimal policies for domains where unforeseen changes occur. The next section presents the action language $BC+$ that provides the appropriate definitions for domain modelling needed to bridge the gap between ASP and RL.

2.3 The Action Language $BC+$

The action language $BC+$ is defined over the stable model semantics and allows for some useful ASP constructs, such as a high-level description of actions and their effects, as a consequence of its structured abstract representation of transition systems [2].

$BC+$ has two sets of symbols: action constants and fluent constants; and also two sets of formulas: fluent formula, which has only fluent constants, and action formula, which has at least one action constant and no fluent constant.

In $BC+$, an action description is a set of causal laws that have two forms. The first is:

$$\text{caused } F \text{ if } G \quad (1)$$

where, F and G are formulas. If F and G are both fluent formulas, then Formula 1 is a static law. If F is an action formula, but G is a fluent formula, then Formula 1 is an action dynamic law. The second form is called fluent dynamic law and has the form:

$$\text{caused } F \text{ if } G \text{ after } H \quad (2)$$

where H is a formula, F and G are fluent formulas and F does not contain statically determined constants.

Causal dependencies between fluents in the same state are described by static laws. Direct effects of actions are represented by fluent dynamic laws, while causal dependencies between concurrently executed actions are expressed by action dynamic laws.

Given an action description \mathcal{D} expressed in $BC+$, a stable model for the sequence $PF_m(\mathcal{D})$ of propositional formulae describes a path of length m in a transition system \mathcal{D} [2]. Given a time instant $i \in \{0, \dots, m\}$, a translation $PF_m(\mathcal{D})$ is a conjunction of:

- $i : F \leftarrow i : G$ for every static and atomic law in \mathcal{D} and $\forall i \in \{0, \dots, m-1\}$;
- $i+1 : F \leftarrow (i+1 : G) \wedge (i : H)$ for every fluent dynamic law in \mathcal{D} and $\forall i \in \{0, \dots, m-1\}$;
- $\{0 : c = v\}$ for every regular fluent constant c and every $v \in \text{Dom}(c)$;

- Given $\{v_1, \dots, v_m\}$ as $\text{Dom}(c)$, $\perp \leftarrow \neg(1 \leq \{i : c = v_1, \dots, i : c = v_m\} \leq 1)$ for every $i : c$ representing the uniqueness of names and existence values for the constants;

The action language $BC+$ can be directly translated into an ASP program for providing sequences of actions as answer sets.

3 Combining ASP and MDP

This section presents the main contribution of this work, the ASP(RL) method, which is a combination of ASP and MDP for solving non-stationary decision making problems.

3.1 Finding the Set of States

In this work Answer Set Programs, translated from $BC+$, represent the states $s \in \mathcal{S}$, the actions $a \in \mathcal{A}$, and the expected transition function of an MDP, along with sets $\mathcal{S}_0 \in \mathcal{S}$ and $\mathcal{S}_g \in \mathcal{S}$ which represents the sets of initial states and goal states respectively. Let $\Pi(\mathcal{S}, \mathcal{A})$ be one such ASP program with \mathcal{S} and \mathcal{A} as set of states and actions respectively. Given an initial state $s_0 \in \mathcal{S}_0$ and a goal state $s_g \in \mathcal{S}_g$, an answer set of $\Pi(\mathcal{S}, \mathcal{A})$ represents a trajectory T of the form:

$$T = \langle \langle s_0, a_0, s_1 \rangle, \langle s_1, a_1, s_2 \rangle, \dots, \langle s_n, a_n, s_g \rangle \rangle \quad (3)$$

where s_n and a_n are, respectively, the state and the action at time n .

As ASP programs can have more than one answer set, let a set \mathcal{H} contain all trajectories T that represent the sequence of actions leading from an initial state to a goal state. Thus, in the set of trajectories \mathcal{H} there are a set of states visited and a set of actions performed that are subsets of those sets in the MDP defined in the logic program $\Pi(\mathcal{S}, \mathcal{A})$. Thus, this set \mathcal{H} can be used to describe a new MDP M , as stated in the following Lemma.

Lemma 1 *Given an MDP $M = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R} \rangle$ described by a logic program $\Pi(\mathcal{S}, \mathcal{A})$, the set \mathcal{H} of trajectories found for $\Pi(\mathcal{S}, \mathcal{A})$ defines $\dot{M} = \langle \dot{\mathcal{S}}, \dot{\mathcal{A}}, \dot{\mathcal{T}}, \dot{\mathcal{R}} \rangle$, such that $\dot{M} \subseteq M$. Considering that $\dot{M} \subseteq M$ iff $\dot{\mathcal{S}} \subseteq \mathcal{S}$ or $\dot{\mathcal{A}} \subseteq \mathcal{A}$ or $\dot{\mathcal{T}} \subseteq \mathcal{T}$ or $\dot{\mathcal{R}} \subseteq \mathcal{R}$.*

Proof (Sketch) A logic program $\Pi(\mathcal{S}, \mathcal{A})$ defines a set of restrictions on an MDP. These restrictions are a set $\bar{\mathcal{S}}$ of states that the agent may not be able to visit and a set $\bar{\mathcal{A}}$ of actions that the agent may not be able to perform. Also, changing actions or states imply changing the transitions as well. Thus, $\dot{\mathcal{S}} \subseteq \mathcal{S} \mid \dot{\mathcal{S}} = \mathcal{S} - \bar{\mathcal{S}}$ and $\dot{\mathcal{A}} \subseteq \mathcal{A} \mid \dot{\mathcal{A}} = \mathcal{A} - \bar{\mathcal{A}}$.

The transition function $\dot{\mathcal{T}}$ is then described considering the following conditions:

1. The agent cannot visit a state that is forbidden: $\bar{\mathcal{S}} \times \mathcal{A}$;

2. The agent cannot perform a forbidden action: $\mathcal{S} \times \overline{\mathcal{A}}$;
3. The agent cannot perform a forbidden action in a state that it cannot visit: $\overline{\mathcal{S}} \times \overline{\mathcal{A}}$;
4. The agent cannot visit states that have no transition probabilities: $\mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto 0$;
5. The agent is not allowed to perform some specific actions in some specific states: $\mathcal{Q}(s, a)$.

Thus, the transition function that is extracted from the answer sets is defined as:

$$\begin{aligned} \dot{\mathcal{T}}(\dot{\mathcal{S}}, \dot{\mathcal{A}}) = & (\dot{\mathcal{S}} \times \dot{\mathcal{A}}) - ((\overline{\mathcal{S}} \times \mathcal{A}) + (\mathcal{S} \times \overline{\mathcal{A}}) + (\overline{\mathcal{S}} \times \overline{\mathcal{A}}) \\ & + (\mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto 0) + (\overline{\mathcal{Q}(s, a)})) \mapsto \dot{\mathcal{S}} \end{aligned} \quad (4)$$

When an MDP is not deterministic, choice rules are used to describe the transition possibilities (without the probability itself), a similar process is used to find the transition function $\dot{\mathcal{T}}$.

Therefore, with this new set of states $\dot{\mathcal{S}} \subseteq \mathcal{S}$, actions $\dot{\mathcal{A}} \subseteq \mathcal{A}$ and transition function $\dot{\mathcal{T}}$, it is possible to formalise an MDP $\dot{M} \subseteq M$ in the form $\dot{M} = \langle \dot{\mathcal{S}}, \dot{\mathcal{A}}, \dot{\mathcal{T}}, \mathcal{R} \rangle$. Since the reward comes from the interaction with the environment, there is no need to suppress any value in this function or even to know which is the reward function beforehand.

Once it is possible to formalize an MDP \dot{M} that is a subset of another MDP M , it is still necessary to guarantee that the optimal solution $\pi_{\dot{M}}^*$ of \dot{M} is the optimal solution of π_M^* of M as stated in Theorem 1.

Theorem 1 *Given a reward function and an evaluation criteria (i.e. maximising or minimising rewards), the optimal solution $\pi_{\dot{M}}^*$ for the MDP $\dot{M} \subseteq M$ is equivalent to the optimal solution π_M^* for the MDP M given the answer sets (trajectories) \mathcal{H} found as solutions to the logic program $\Pi(\mathcal{S}, \mathcal{A})$ that represents M .*

Proof Both M and \dot{M} have to maximise (or minimise) the same reward function \mathcal{R} . If there is no restrictions in the set of states ($\overline{\mathcal{S}} = \emptyset$) and actions ($\overline{\mathcal{A}} = \emptyset$), we have that $\dot{M} = M$ and $\pi_M^* = \pi_{\dot{M}}^*$.

If there are restrictions represented in $\Pi(\mathcal{S}, \mathcal{A})$, then $\dot{M} \subset M$ and the feasible solutions (answer sets) \mathcal{H} for M are the same of those for \dot{M} (by using Lemma 1). Since the optimal solution must be a feasible solution, then $\pi_M^* \in \mathcal{H}$ and $\pi_{\dot{M}}^* \in \mathcal{H}$. Thus, given the same set of feasible solutions and the same evaluation criteria, $\pi_M^* = \pi_{\dot{M}}^*$.

3.2 The Algorithm ASP(RL)

Lemma 1 and Theorem 1 support the use of ASP to find the sets of states and actions of an MDP. By using RL it is possible to find an optimal stochastic

```

1 Algorithm: ASP(RL)
   Input: An MDP described as a logic program  $\Pi(\mathcal{S}, \mathcal{A})$  and a (optional)  $Q(s, a)$ 
           function to be approximated.
   Output: The approximated  $Q(s, a)$  function.
2 Find the answer sets  $\mathcal{H}$  for  $\Pi(\mathcal{S}, \mathcal{A})$ .
3 Update  $Q(s, a)$  function using  $\dot{\mathcal{S}}$  and  $\dot{\mathcal{A}}$  found in  $\mathcal{H}$ .
4 while the environment does not change do
5   | Approximate  $Q(s, a)$  using a RL method.
6 end
7 Include the observed changes in  $\Pi(\mathcal{S}, \mathcal{A})$ 
8 Call ASP(RL) with  $\Pi(\mathcal{S}, \mathcal{A})$  and the  $Q(s, a)$  function approximated.

```

Algorithm 1: ASP(RL) Algorithm.

solution to this MDP. Since ASP allows for revisions to be made in the set of states and actions, if it is the case that the environment changes at any time step, it can be used to find the new subsets $\dot{\mathcal{S}}$ of states and $\dot{\mathcal{A}}$ of actions of the modified MDP and values learnt from the previous interaction can be used as input for this new MDP. Algorithm 1 is the pseudocode of ASP(RL), that uses the non-monotonicity of ASP along with the exploratory nature of RL algorithms in stochastic domains.

Algorithm 1 uses RL methods for approximating the $Q(s, a)$ function for the states and actions obtained by ASP. First, the domain is described as a logic program $\Pi(\mathcal{S}, \mathcal{A})$, using the *BC+* vocabulary, and answer sets are found for it. From those answer sets (as shown in Lemma 1) the sets of states $\dot{\mathcal{S}}$ and actions $\dot{\mathcal{A}}$ are constructed for the MDP that will be used by the agent to interact with the environment, along with the transition function $\dot{\mathcal{T}}$. Once the MDP $\dot{M} = \langle \dot{\mathcal{S}}, \dot{\mathcal{A}}, \dot{\mathcal{T}}, \mathcal{R} \rangle$ is formalised, the interaction with the environment and the search for the optimal solution begins by using any RL algorithm. This interaction continues until a change in the environment happens. At this instant, the algorithm returns the approximated $Q(s, a)$.

The algorithm works in non-stationary environments by including the observed environment changes in $\Pi(\mathcal{S}, \mathcal{A})$ so that ASP can be used again to find the new sets of states and actions along with the transition function. Since there is a $Q(s, a)$ function approximated from the previous interaction, modifications are performed in it. The state-action pairs that are in the new set of answer sets are added to the action-value function and the pairs that are not in this set are removed. The state-action pairs that were in the function, and that are also in the answer set, remain in the action-value function with the previously learned value. Therefore, the interaction with a changing environment is done by calling ASP(RL) with $\Pi(\mathcal{S}, \mathcal{A})$, augmented with the observed changes, and the action-value function returned by the previous call.

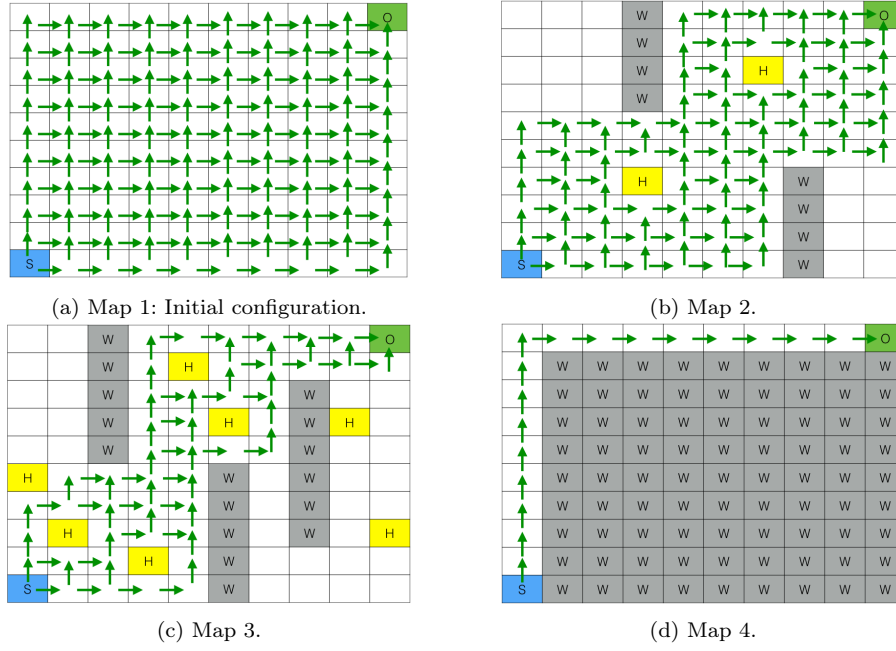


Fig. 1: Grid worlds used in the experiments. Squares labelled with W represent walls and with H, holes.

4 Experiments

Experiments were performed in a non-deterministic non-stationary grid world of size 10×10 which allowed the execution of only one of four actions each time: *go up*, *go down*, *go left* and *go right*. The probabilities for the environment were defined as 80% for the transition to happen as expected (e.g., executing *go up* makes the agent go up with 80% of probability) and a 20% chance for the agent to go orthogonal to the desired direction (e.g., executing *go up* may make the agent to go left or right with 10% of chance for each side).

The grid world may have walls (W) and holes (H) each of which occupies a single cell of the grid. When the agent performs an action and hits a wall, it stays in the same state; when it executes an action and falls into a hole, the episode ends. In this domain, an agent that starts in the lowermost, leftmost, cell has as a goal to reach the topmost, rightmost, cell. The reward function used in this domain is +100 for reaching the goal, -100 for falling in a hole and -1 in any other event. It is important to notice that the transition function and reward function are unknown to the agent. For this grid world, the representation used by the agent is the value of its position in X and Y. These values are not treated by the agent as an X by Y matrix, but as a set of atoms in the form (X, Y) for each pair of X and Y values found in an answer set.

This grid world suffers changes in a manner that is previously unknown to the learning agent. In this work, ASP(RL) is evaluated in three distinct situations, in each of them the agent starts in the map shown in Figure 1a that, after 5000 episodes, changes to one of the other maps in Figure 1. For this work, changes observed in the environment were manually entered in the logic program. Nevertheless, this can be automatically done by using an online method with ASP.

The map in Figure 1a represents a grid world with no walls or holes. In this case, any combination of actions that makes the agent to go *up* and *right* leads the agent to the goal. Figure 1b represents a grid world with two walls and two holes. Figure 1c shows a grid world containing more walls and holes than in the previous situation. In this case, the agent has fewer action options to achieve the goal state. Finally, Figure 1d represents a grid world in which there is only one policy for achieving the goal state with the minimum number of actions. Any other policy for this grid world will necessarily make the agent hit a wall before reaching the goal state.

The arrows in the maps shown in Figure 1 represent the feasible policies obtained by ASP with the minimum number of steps. Note that, these policies do not represent the transition probabilities of the environment.

In the first situation, the environment changes from the map in Figure 1a to that in Figure 1b. In this case, we can see that there is a reduction in the number of policies with the minimum number of steps.

In the second situation, the change occurs from the map shown in Figure 1a to that in Figure 1c. By analysing the arrows in the final (Fig. 1c) grid world, we can see that there is an even greater reduction in the number of policies than in the previous situation, since there are more walls and holes in this map, which imply fewer safe actions (*arrows*) available.

In the final situation the environment changes from the map in Figure 1a to the one in Figure 1d. In this case the MDP has only one optimal solution. This situation was chosen since the answer set provides the only optimal solution almost instantly, whereas in the case where an action-value function is approximated by RL (without using the answer sets), every possible action in every possible state is considered, leading to a costly search procedure.

5 Results

In this section we use the situations described above to compare the learning processes of SARSA and Q-Learning with those of ASP(SARSA) and ASP(Q-Learning), which are ASP(RL) methods where SARSA and Q-Learning are used along with ASP. This comparison is accomplished with two different criteria: the return ($\sum r(s, a, s')$) of the episode and the number of steps needed to reach the goal state and root-mean-square deviation (RMSD) of the action-value function, at time t wrt time $t - 1$, according to the equation 5 below.

$$\text{RMSD} = \sqrt{\frac{\sum_{m=1}^n Q_{m,t}(s, a) - Q_{m,t-1}(s, a)}{n}} \quad (5)$$

The graphs in Figures 2, 3, 4, 5, 6 and 7 represent measurements for the four algorithms applied in the three situations considered. Figures 3, 5 and 7 depict the results of the first 50 episodes for the first map and then skipping to the 5000th episode directly, in order to present the measurements after the environment change occurs¹. The respective results for every episode are shown in Figures 2, 4 and 6.


The results for the number of steps from the first situation are presented in Figure 3a. Figure 3b presents the returns. In the first map (Figure 1a) all four algorithms present the same number of steps and returns during the initial 50 episodes shown. After the 5000th episode, the number of steps of ASP(Q-Learning) and ASP(SARSA) decrease faster than that of Q-Learning and SARSA, while the returns of ASP(Q-Learning) and ASP(SARSA) increase faster than Q-Learning and SARSA. This difference in the performance of ASP(RL) and RL algorithms after the change occurs in the environment is due to the fact that ASP(RL) reuses the $Q(s, a)$ function approximated in the previous map.

For the second situation, Figure 5a presents the number of steps and Figure 5b the returns for the four algorithms. Regarding the number of steps, it is possible to notice that although ASP(Q-Learning) and ASP(SARSA) use the information acquired from previous experience, they still need the same number of steps as Q-Learning and SARSA in all episodes. However, the returns for ASP(Q-Learning) and ASP(SARSA) are higher than the returns from Q-Learning and SARSA when the change in the map occurs (5000th episode). This similarity in the number of steps for the four algorithms is due to the great change that occurred in the environment, thus ASP(Q-Learning) and ASP(SARSA) still need to learn interactively with the new environment, even though they use information from the previous map.

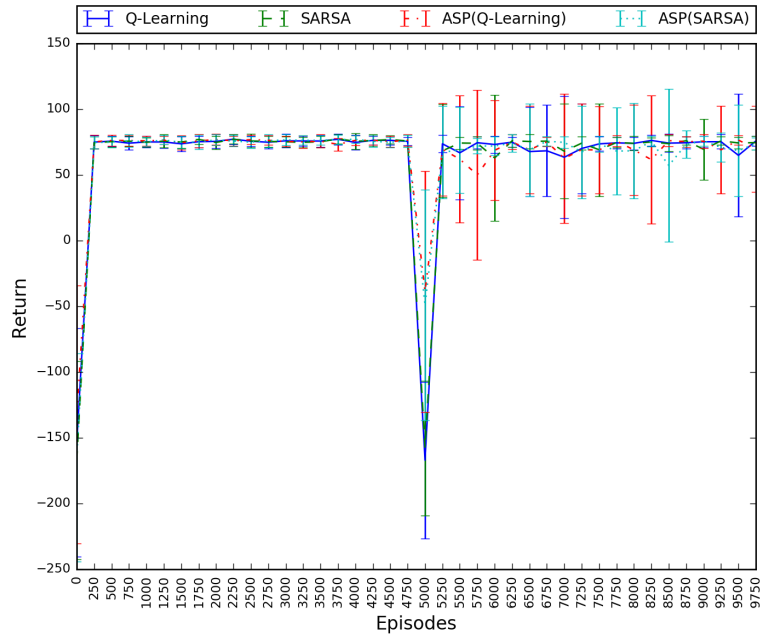
The number of steps and returns for the third situation are presented in Figure 7a and 7b respectively. In both returns and steps, when the change occurs, the use of previously learned values enhance the performance of ASP(Q-Learning) and ASP(SARSA). While there is a slow decrease in the number of steps in Q-Learning and SARSA and a slow increase in the returns, ASP(Q-Learning) and ASP(SARSA) can quickly learn the only optimal policy since this policy is already known from previous map.

Experiments were performed in a 1.66GHz Core2Duo with 4GB of RAM running Debian 9 (currently the testing version). Logic programs were written in *BC+* [2] and translated to ASP language using *CPLUS2ASP* [1], which uses *iClingo* [12] to find answer sets. For finding the optimal solution, Q-Learning and SARSA were implemented in Python 3.5 using only built-in libraries. Thirty training sessions were executed for each algorithm. The same parameters were used in all the experiments: learning rate $\alpha = 0.2$, discount factor $\gamma = 0.9$, exploration/ exploitation rate $\epsilon = 0.1$ and the Q table was randomly initialised.

¹ Episodes 51st to 4999th were removed from the figures.

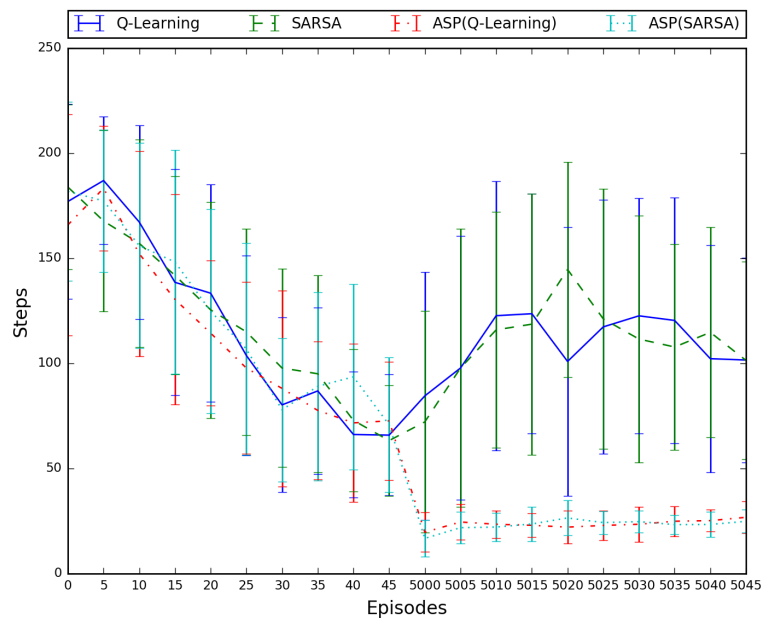


(a) Steps needed to reach the goal state.

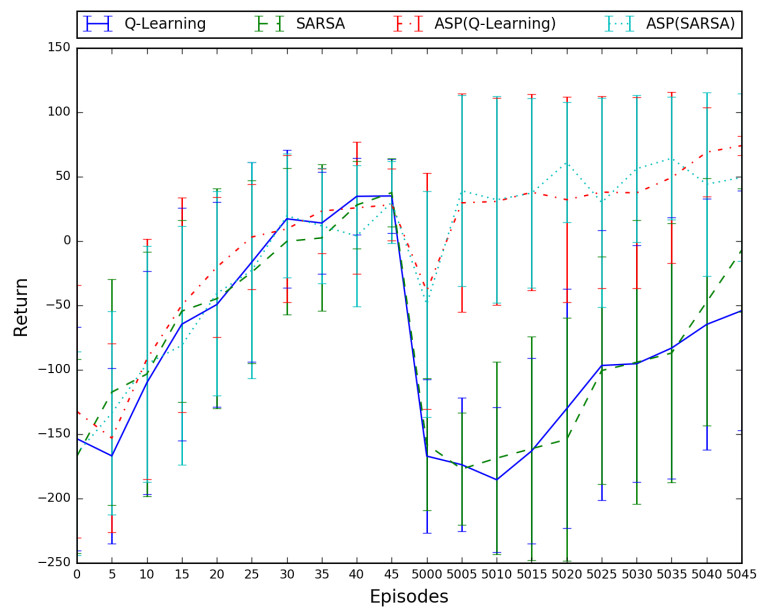


(b) Total returns receive per episode.

Fig. 2: Results for the first situation for every episode.




(a) Steps needed to reach the goal state.




(b) Total returns received per episode.

Fig. 3: Results for the first situation.




(a) Steps needed to reach the goal state.




(b) Total returns receive per episode.

Fig. 4: Results for the second situation for every episode.

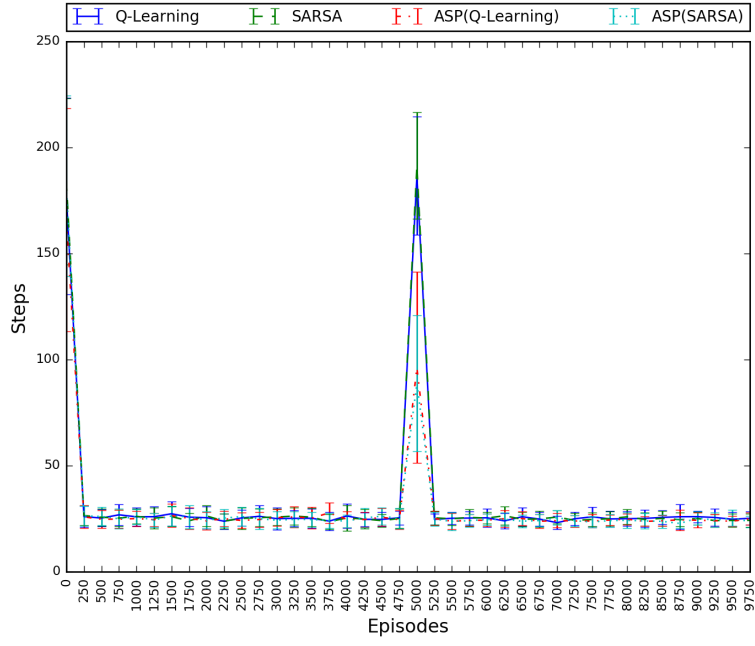


(a) Steps needed to reach the goal state.




(b) Total returns receive per episode.

Fig. 5: Results for the second situation.

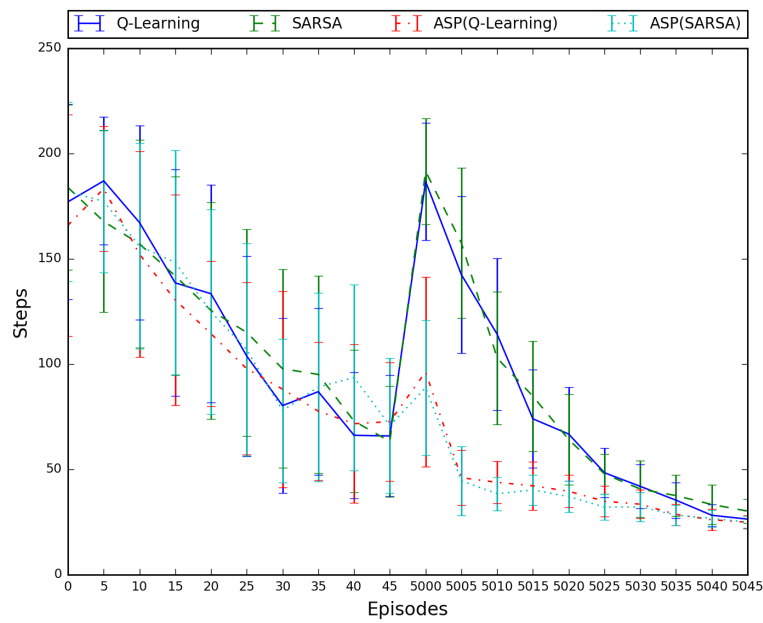


(a) Steps needed to reach the goal state.




(b) Total returns receive per episode.

Fig. 6: Results for the third situation for every episode.



(a) Steps needed to reach the goal state.



(b) Total returns receive per episode.

Fig. 7: Results for the third situation.

6 Discussion

The results shown in the previous section, present the best, worst and average cases of the ASP(RL) method proposed.

The first map (Figure 1a) represents the worst case for ASP(RL). As can be seen in the graphs in Figures 2, 3, 4, 5, 6,7, the performance of ASP(Q-Learning) and ASP(SARSA) are the same as that of Q-Learning and SARSA. This is due to the fact that the reduction in the sets of states and actions are minimal (since there isn't any restriction in this map) and ASP(RL) methods use the same \mathcal{S} and \mathcal{A} as an RL method.

The best case is represented in the last map (Figure 1d). In this case, there is only one feasible policy and, thus, this is the optimal policy. Although the learning process has been executed, in situations like these learning is not necessary, since there is only one feasible policy that is provided by an answer set.

A similar case occurs when there is no feasible policy. In this situations there is also no need to perform the learning process, since it is already known from the answer sets that there is no feasible/optimal policy and the problem cannot be solved.

The average case is presented in the second and third maps (Figures 1b and 1c respectively). In these situation it is possible to notice that there is a reduction in the sets of states and actions, along with a reduction in the search space. Nevertheless, the acceleration in the learning process depends on how much the environment has changed from the previous situation. For example, the gain in learning time in the second situation (Figures 4 and 5) is greater than that of the third situation (Figures 6 and 7).

ASP(RL) was not only capable of dealing with non-stationary non-deterministic environments but it also provides the possibility to reduce the search space, thus finding the optimal solution in fewer interactions with the environment than using RL alone. This reduction in the search space is related to the problem that is being solved and not only to the method proposed.

7 Related Work

The method proposed in this paper is in line with the work reported in [27,22] where ASP is used to find a description of the domain and RL is applied in the search for the optimal solution. Although both proposals combine similar tools, their use differ. While the present work formalises an MDP from the answer sets, the method proposed in [27,22] finds only one answer set for the problem, where each atom in this set defines a hierarchical POMDP that has to be solved.

A related approach is the combination of ASP with action costs [15,25]. Although this method also uses a logic program to describe the domain, it uses a method different from RL to find the action costs. At each action that is executed, the agent finds new plans to reach the goal; the update of the

state-action pair's value is not based on the Temporal Difference method of RL.

Another work that also deals with sequential decision making is P-Log [14, 6] which calculates transition probabilities from sampling the environment, but without considering the cost of performing an action. The present work differs from P-Log in that our goal is to find the optimal solution regarding not only the transition probabilities, but also the action costs.

Also related to our work is Saturated Path-Constrained MDP (SPC-MDP) [21]. In a SPC-MDP, a solution is found by a constraint satisfaction procedure. This closely relates to the results obtained with the use of ASP to define the set of states for an MDP as proposed in this paper. However, while the approach described in [21] uses a Dynamic Programming algorithm to find the solutions, ASP(RL) uses the interaction with the environment in order to approximate the action-value function in non-stationary decision making problems, which (to the best of our knowledge) has never been attempted before.

Works that are somewhat related to our approach, but can be used when searching for the optimal policy are the ones that deals with changing reward functions, such as [10, 11, 26]. Since ASP(RL) uses RL, changes in the reward function are learned by the agent and does not affect the algorithm. Another approach that is somewhat related to ASP(RL) is hierarchical MDPs (such as the works of [19, 20]), which can also be incorporated such as the method proposed by [27, 22] described in the beginning of this section. Although the decomposition proposed by hierarchical MDPs provide more abstraction when searching for the solution, ASP(RL) deals with changes in \mathcal{S} and \mathcal{A} such as the number of states and actions available in the environment or their representation.

To the best of our knowledge, these are the only works related to our method in which the focus is in the change of the sets of states and actions, not only in the transition and reward functions, nevertheless comparison with these methods is not possible since their goals and results differs from ASP(RL).

8 Conclusion

This paper presented a method for efficiently solving non-stationary Markov Decision Processes (MDP). The proposed approach, called ASP(RL), uses a combination of Answer Set Programming (ASP) and Reinforcement Learning (RL) in which ASP provides the set of states and actions in domains where unforeseen changes may happen, while RL is used to approximate a value-action function by means of interactions with the environment. In ASP(RL), Answer Set Programming is used as a tool for reasoning and knowledge revision and Reinforcement Learning allows for learning the solution of an MDP without the need of an explicit stationary reward function.

Experiments were performed in a changing grid world, whose results show that the use of ASP to find the set of states and actions effectively reduces the search space for finding optimal policies of Markov Decision Processes

in complex domains, as well as in domains that allow only a few possible policies. Not only ASP(RL) allowed a faster approximation of the action-value function (compared to standard RL algorithms), but the process could continue to interact in a changing environment indefinitely.

ASP(RL) is capable of dealing with unforeseen changes in the domain, thus solving non-stationary decision making problems. To the best of our knowledge, this has never been accomplished before.

Future work shall be directed towards a full integration of RL into the ASP engine, facilitating the use of ASP when new states appear in a non-deterministic environment, with the possibility of reviewing the whole set of states seamlessly.

References

1. Babb, J., Lee, J.: Cplus 2ASP: Computing action language $\mathcal{C} +$ in answer set programming. In: P. Cabalar, T.C. Son (eds.) *Logic Programming and Nonmonotonic Reasoning*, vol. 8148, pp. 122–134. Springer Berlin Heidelberg (2013)
2. Babb, J., Lee, J.: Action language $BC+$. *Journal of Logic and Computation* (2015)
3. Balduccini, M., Gelfond, M., Nogueira, M.: Answer set based design of knowledge systems. *Annals of Mathematics and Artificial Intelligence* **47**(1-2), 183–219 (2006)
4. Balduccini, M., Gelfond, M., Nogueira, M., Watson, R.: Planning with the USA-Advisor. In: 3rd NASA International workshop on Planning and Scheduling for Space. Houston, Texas (2002)
5. Balduccini, M., Gelfond, M., Watson, R., Nogueira, M.: The USA-Advisor. In: G. Goos, J. Harmanis, J. van Leeuwen, T. Eiter, W. Faber, M.I. Truszczyński (eds.) *Logic Programming and Nonmonotonic Reasoning*, vol. 2173, pp. 439–442. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
6. Baral, C., Gelfond, M., Rushton, N.: Probabilistic reasoning with answer sets. *Theory and Practice of Logic Programming* **9**(1), 57 (2009)
7. Bellman, R.: On the theory of dynamic programming. *Proceedings of the National Academy of Sciences* **38**(8), 716–719 (1952)
8. Bellman, R.: A Markovian decision process. *Indiana University Mathematics Journal* **6**(4), 679–684 (1957)
9. Bellman, R.E., Dreyfus, S.E.: *Applied dynamic programming*, 4 edn. Princeton Univ. Press (1971)
10. Even-dar, E., Kakade, S.M., Mansour, Y.: Experts in a markov decision process. In: L.K. Saul, Y. Weiss, L. Bottou (eds.) *Advances in Neural Information Processing Systems* 17, pp. 401–408. MIT Press (2005)
11. Even-Dar, E., Kakade, S.M., Mansour, Y.: Online markov decision processes. *Mathematics of Operations Research* **34**(3), 726–736 (2009)
12. Gebser, M., Kaminski, R., Kaufmann, B.: *Answer set solving in practice*. Morgan & Claypool Publishers (2013)
13. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In: R. Kowalski, Bowen, Kenneth (eds.) *Proceedings of International Logic Programming Conference and Symposium*, pp. 1070–1080. MIT Press (1988)
14. Gelfond, M., Rushton, N.: Causal and probabilistic reasoning in P-log. *Heuristics, Probabilities and Causality. A tribute to Judea Pearl* pp. 337–359 (2010)
15. Khandelwal, P., Yang, F., Leonetti, M., Lifschitz, V., Stone, P.: Planning in action language BC while learning action costs for mobile robots. In: *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS 2014, Portsmouth, New Hampshire, USA, June 21-26, 2014* (2014)
16. Lifschitz, V.: Answer set programming and plan generation. *Artificial Intelligence* **138**(1-2), 39 – 54 (2002)

17. McCarthy, J.: Elaboration tolerance. In: Proc. of the Fourth Symposium on Logical Formalizations of Commonsense Reasoning (Common Sense 98), vol. 98. London, UK (1998)
18. Nogueira, M., Balduccini, M., Gelfond, M., Watson, R., Barry, M.: An A-Prolog decision support system for the space shuttle. In: G. Goos, J. Hartmanis, J. van Leeuwen, I.V. Ramakrishnan (eds.) *Practical Aspects of Declarative Languages*, vol. 1990, pp. 169–183. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
19. Ryan, M.R., Pendrith, M.D.: RL-tops: An architecture for modularity and re-use in reinforcement learning. In: *Proceedings of the Fifteenth International Conference on Machine Learning*, pp. 481–487. Morgan Kaufmann (1998)
20. Ryan, M.R.K.: Using abstract models of behaviours to automatically generate reinforcement learning hierarchies. In: *Proceedings of The 19th International Conference on Machine Learning*, pp. 522–529. Morgan Kaufmann (2002)
21. Sprael, J., Teichteil-Königsbuch, F., Kolobov, A.: Saturated path-constrained MDP: Planning under uncertainty and deterministic model-checking constraints. In: Proc. of 28th AAAI Conf. on Artificial Intelligence (AAAI), pp. 2367–2373 (2014)
22. Sridharan, M., Gelfond, M., Zhang, S., Wyatt, J.: Mixing non-monotonic logical reasoning and probabilistic planning for robots. In: *Workshop on Hybrid Reasoning @ IJCAI 2015* (2015)
23. Sutton, R.S., Barto, A.G.: *Reinforcement learning an introduction – Second edition*, in progress (Draft). MIT Press (2015)
24. Watkins, C.J.C.H.: *Learning from delayed rewards*. PhD thesis, University of Cambridge England (1989)
25. Yang, F., Khandelwal, P., Leonetti, M., Stone, P.: Planning in answer set programming while learning action costs for mobile robots. In: *AAAI Spring 2014 Symposium on Knowledge Representation and Reasoning in Robotics (AAAI-SSS)* (2014)
26. Yu, J.Y., Mannor, S., Shimkin, N.: Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research* **34**(3), 737–757 (2009)
27. Zhang, S., Sridharan, M., Wyatt, J.L.: Mixed logical inference and probabilistic planning for robots in unreliable worlds. *IEEE Transactions on Robotics* **31**(3), 699–713 (2015)