

# FROM REACTIVE TO COGNITIVE AGENTS: EXTENDING REINFORCEMENT LEARNING TO GENERATE SYMBOLIC KNOWLEDGE BASES

RÔMULO GUEDES CERQUEIRA\* AUGUSTO LOUREIRO DA COSTA\*  
STEPHEN G. MCGILL† DANIEL LEE† GEORGE PAPPAS†

*\*Robotics Laboratory,  
Postgraduate Program in Electrical Engineering  
Federal University of Bahia, 40210-630 Salvador, Bahia, Brazil*

*†GRAP Lab,  
University of Pennsylvania, Philadelphia, Pennsylvania, USA*

Email: romulogcerqueira@gmail.com, augusto.loureiro@ufba.br,  
smcgill13@seas.upenn.edu, lee@seas.upenn.edu, pappasg@seas.upenn.edu

**Abstract**— A new methodology for knowledge-based agents to learn from interactions with their environment is presented in this paper. This approach combines Reinforcement Learning and Knowledge-Based Systems. A Q-Learning algorithm obtains the optimal policy, which is automatically coded into a symbolic rule base, using first-order logic as knowledge representation formalism. The knowledge base was embedded in an omnidirectional mobile robot, making it able to navigate autonomously in unpredictable environments with known and unknown obstacles using the same knowledge base. Additionally, a method of space abstraction based in human reasoning was formalized to reduce the number of complex environment states and to accelerate the learning. The experimental results of autonomous navigation executed by the real robot are also presented here.

**Keywords**— Robot Navigation, Knowledge-based Systems, Reinforcement Learning, Embedded Intelligent Systems.

**Resumo**— Uma nova metodologia para agentes baseados em conhecimento aprenderem através de interações com seus ambientes é apresentada neste artigo. Esta abordagem combina Aprendizado por Reforço e Sistemas Baseados em Conhecimento. Um algoritmo Q-Learning obtém a política ótima, que é automaticamente codificada para uma base de regra simbólica, utilizando lógica de primeira ordem como formalismo de representação do conhecimento. A base de conhecimento foi embarcado em um robô móvel omnidirecional, tornando-o capaz de navegar autonomamente em ambientes imprevisíveis com obstáculos conhecidos e desconhecidos usando a mesma base de conhecimento. Adicionalmente, um método de abstração do espaço baseado em raciocínio humano foi formalizado para reduzir o complexo número de estados do ambiente e acelerar o aprendizado. Os resultados experimentais da navegação autônoma executada pelo robô real também é apresentado aqui.

**Palavras-chave**— Navegação de Robôs, Sistemas baseados em Conhecimento, Aprendizado por Reforço, Sistemas Inteligentes Embarcados.

## 1 Introduction

In recent years, the academic community of Artificial Intelligence has searched to replace explicit programming by the process of teaching a task. Autonomous robot skill acquisitions have presented a large amount of works using learning processes. In this context, Reinforcement Learning (RL) is a widely used technique, where an agent learns a policy  $\pi$ , exploring a partially or totally unknown environment, based in rewards obtained from its actions. Thus, the optimal policy  $\pi^*$  maps each environment state  $s_t$  to the best action  $a_t$  and it is stored in a hash table.

In task execution, this method presents the classical reactive agents limitations, some situation, like when the robot is surrounded by mobile obstacles, it collapses. Besides, the machine learning process usually takes a long time to learn a new desired behaviour. New learning episodes also are needed for each new environment configurations. In unpredictable changing environments with unpredictable obstacles, the agent can spend more time learning than executing the de-

sired task. Moreover, the storage and management of several look-up tables becomes hard work for a reactive agent.

A two layered approach, where basic behaviours are implemented at a low level layer and supervisory fuzzy rule based system is implemented at a high level layer, decides the appropriate basic behaviours in the current environment state is presented in (Fatmi et al., 2006). The fuzzy rules base was automatically generated by Reinforcement Learning algorithm. This improves the behaviour based navigation idea in (Brooks, 1991), and extends the Reinforcement Learning to generates a fuzzy rules based system. In another approach (Torrey, 2009), Inductive Logic Programming is used to transfer learning, where relevant knowledge from previous learning experiences was applied to help the new task learning. Recently, a significant amount of interesting research works about mobile robot navigation has been presented. In these works, some widely utilized methods are neural networks in (Ouarda, 2010), fuzzy logic in (Shayestegan and Marhaban, 2012), reinforce-

ment learning in (Jaradat et al., 2011), genetic algorithms in (Manikas, 2007), and rule-based-neuro-fuzzy technique in (Pradhan et al., 2007).

Looking forward to give full decision autonomy to mobile robots, in high level tasks executions, a cognitive agent called Concurrent Autonomous Agent architecture (CAA) (da Costa and Bittencourt, 1999), was chosen to be embedded in the respective mobile robot. This agent architecture, embodied a model for higher-level cognition that supports learning and reasoning mechanisms using logical expressions. The cognitive model is based on the following fundamental hypotheses (Bittencourt, 2007):

- Cognition is an emergent property of a cyclic dynamic self-organizing process based on the interaction of a large number of functionally independent units of a few types.
- Any model of the cognitive activity should be epistemologically compatible with the Theory of Evolution (Mccarthy and Hayes, 1969). That applies not only to the “hardware” components of this activity but also to its “psychological” aspects.
- Learning and cognitive activity are closely related and, therefore, the cognitive modelling process should strongly depend on the cognitive agents particular history (Bittencourt, 2007).

Since learning and cognitive activity are closely related, the methodology that combines RL and Knowledge-Based Systems, was implemented to allow the CAA learns with its interaction with the environment. This approach makes CAA fully compatible with the three hypothesis for a cognitive model. In the previous CAA implementation, the learning propriety would be eligible just for reactive level, since that both instinctive and cognitive level uses Knowledge Based Systems for automatic reasoning. The methodology was applied to the instinctive level, where a knowledge-based system is used to plan execution. For a given goal  $g_i$ , the RL obtains the optimal policy which is automatically coded using a knowledge representation formalism based on first order logic. Each knowledge base  $KB_i$  stores the necessary knowledge for the agent to achieve the goal  $g_i$ , combining a set of basic behaviours, that can be performed in the environment. A new approach for environment mapping around the mobile robot neighbourhood allows a significant reduction of the relevant state space.

The paper is organized as follows: In Section 2, the proposed approach for abstraction of mobile robot neighbourhood is presented. Section 3 describes the algorithm based on RL, where the optimal policy  $\pi^*$  is learned, and then coded into

the knowledge bases  $KB_i$ . Section 4, presents experimental results from simulated mobile robot locomotion, and real mobile robot locomotion. Section 5 discusses conclusions and future works.

## 2 Reduced States Space Mapping

When a human tries to walk safely in a dynamic environment, he will care for the relative or approximate distances and directions between him and the goal and the closest obstacle, navigating safely independently of the number of obstacles. In this work, the agent learns based on human reasoning. The robot was considered to be the center of universe and the environment surrounding was divided into four orthogonal regions: R1, R2, R3 and R4, as seen in figure 1.




Figure 1: The Environment Zones and States Space

The environment state at each time  $t$  is determined by the goal region ( $R_g$ ) and closest obstacle region ( $R_o$ ):

$$S_t = (R_g, R_o), t \in [1, 16] \quad (1)$$

This approach dramatically reduces the environment to 16 states, accelerates Q-Learning, decreases computational cost, and more specifically generates a small knowledge base, which is very appropriate for embedding in mobile robots.

This work uses global vision system, where the environment surrounding the robot is known, but could also use an omnidirectional local vision system. This works adopts some assumptions:

1. The robot’s position and velocity are known at each instant;
2. The goal’s position and velocity are known at each instant;
3. The closest obstacle’s position and velocity are known at each instant.

The new methodology consists of three steps: agent learning, simulated environment execution, and embedded knowledge base and task execution into mobile robot. The first and second steps are done into a simulated environment, using Player/Stage robotics simulator. Once the optimal policy  $\pi^*$  is generated it is coded into a symbolic knowledge Base  $KB$  which is embedded into CAA’s instinctive level. Then new simulation

experiments, where the environment assumes different configurations, are performed. Finally the agent is embedded into the omnidirectional mobile robot, and the free collision navigation experiments are done. In a second experiment, now using a NaO humanoid robot, the agent reactive level was replaced by another version which embedded a Zero Moment Point Controller and the local Vision systems, provided by the Unified Humanoids Robotic Software Platform (McGill et al., 2010), and that instinctive level and mainly the Knowledge Bases was kept the same, and the same free collision navigation experiments were successful repeated with the humanoid robot.

### 3 Reinforcement Learning for Trajectory Planning

The RL model that CAA uses in this work, is presented in figure 2. At each time instant  $t$  (iteration), the agent uses its sensors to perceive the current environment's state  $s_t$ . Following a policy  $\pi$ , an action  $a_t$  is chosen and executed by its actuators. This action leads the environment to assume a new state  $s_{t+1}$ . The agent is rewarded if its taken action was good; otherwise it is punished. This reinforcement measurement, choice of the two works over time, will cause a correction of policy, defining a dynamic, interactive and learning process.

A  $\pi$  policy maps the environment states into actions, and defines the agent behaviour over time. Thus, the role of RL system is to find an optimal policy  $\pi^*$  that maximizes the reward  $R$  returned by environment. Once the optimal policy  $\pi^*$  is learned, it is coded in knowledge bases.




Figure 2: RL model with CAA architecture

The proposed algorithm, implemented in C programming language, is based in Q-Learning, a machine learning technique that has been successfully utilized for solving the mobile robot navigation problem in dynamic environments.

#### 3.1 The Set of Actions

An omnidirectional robot and robot soccer field, with Robocup F180 League official dimensions, composes the environment in this work. According to the holonomic constraints of an omnidirectional robot, the set of actions is defined

by eight movement actions:

*North, Northeast, East, Southeast, South, Southwest, West, Northwest.*

#### 3.2 The Reward Function

The reward function is defined by the euclidean distances between the robot and goal ( $d_{rg}$ ) and closest obstacle ( $d_{ro}$ ). Thus, the set of states is classified as follows:

- Safe Zone ( $SZ$ ), when the robot has no or low possibility of collision with any obstacle;
- Not-Safe Zone ( $NZ$ ), when the robot has high possibility of collision with any obstacle;
- Victory Zone ( $VZ$ ), when the robot reaches its goal (the ball in this case);
- Failure Zone ( $FZ$ ), when the robot collides with an obstacle.

Moreover:

- $rad_{vz}$  is the radius of  $VZ$  around the goal;
- $rad_{nz}$  is the radius of  $NZ$  around the closest obstacle;
- $rad_{col}$  is the radius around closest obstacle.

The robot zone at time instant  $t$  (figure 1), is defined as:

$$Z(t) = \begin{cases} VZ, d_{rg} \leq rad_{vz} \\ FZ, d_{ro} \leq rad_{col} \\ SZ, d_{ro} > rad_{nz} \\ NZ, rad_{nz} < d_{ro} \leq rad_{col} \end{cases} \quad (2)$$

From data, the reward function is defined:

$$r = \begin{cases} 4, Z \subset SZ \rightarrow VZ \\ 4, Z \subset NZ \rightarrow VZ \\ 1, Z \subset NZ \rightarrow SZ \\ -1, Z \subset SZ \rightarrow NZ \\ 3, Z \subset NZ \rightarrow NZ, d_{ro}(n+1) < d_{ro}(n), \\ \quad d_{rg}(n+1) < d_{rg}(n), d_{ro}(n+1) \geq d_{rb}(n+1) \\ -3, Z \subset NZ \rightarrow NZ, d_{ro}(n+1) < d_{ro}(n), \\ \quad d_{rg}(n+1) < d_{rg}(n), d_{ro}(n+1) < d_{rb}(n+1) \\ -2, Z \subset NZ \rightarrow NZ, d_{ro}(n+1) < d_{ro}(n), \\ \quad d_{rg}(n+1) \geq d_{rg}(n) \\ 2, Z \subset NZ \rightarrow NZ, d_{ro}(n+1) \geq d_{ro}(n), \\ \quad d_{rg}(n+1) < d_{rg}(n) \\ -1, Z \subset NZ \rightarrow NZ, d_{ro}(n+1) \geq d_{ro}(n), \\ \quad d_{rg}(n+1) \geq d_{rg}(n) \\ 0, Z \subset SZ \rightarrow SZ \\ -4, Z \subset NZ \rightarrow FZ \end{cases} \quad (3)$$

#### 3.3 The Value Function

Major RL algorithms are based on value functions. Q-Learning stores the value function in a Q-table, initially null. During the learning process, Q-table is filled by seeing different scenarios. This table

has an immediate reward value from the taken action and the maximum accumulated reward for the possible actions in a current state (see Equation 4).

$$Q(s_t, a_t) = r(s_t, a_t) + \gamma * \max(Q, s_{t+1}) \quad (4)$$

where:  $s_t$  is the state at time instant  $t$ ;  $a_t$  is the action taken at  $t$ ;  $r(s_t, a_t)$  is the immediate reward received when an action  $a_t$  in a state  $s_t$  is taken;  $\gamma$  is the reduction factor, a number in the range of [0,1] that determines if the learned policy will consider delayed or immediate rewards, and  $\max(Q, s_{t+1})$  is the maximum Q-value calculated for taking all possible actions on the state at the next time step.

### 3.4 Learning Process

The learning process is composed by a given number of episodes. For each episode, one give scenario is presented by the agent that chooses an action, considering the goal and the expected reward. Q-table is updated by equation 4. Each scenario consists of the following: the robot position, the desired position, and just one static obstacle. Since the robot only considers with the nearest obstacle, then one obstacle is enough for training.

The element positions are chosen randomly for each episode. The robot always moves toward the goal, independent of a possible collision. It allows the Q-table to converge to the optimal policy  $\pi^*$ , and the robot remains in SZ until it reaches the desired destination.

Once inside SZ, the agent move towards the goal. If the last action lead to a state transition to a VZ or CZ, the episode ends. Thus, the Q-table is updated only when the agent is inside NZ. If it is the last episode, learning is finalized. The learning process is presented in figure 3.

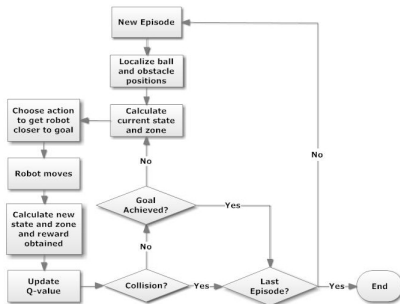


Figure 3: Learning Process

At the end of learning, an algorithm written in C language converts automatically the optimal policy presented by Q-table into a rule base  $KB$ . Given the set of environment states  $S$  and the set of possible actions  $A$ , the biggest Q-value for each state presents the best action to be executed by

agent. For these experiments, six rules base  $KB$  where generated, which contains the optimal policies, for 10,000, 25,000, 50,000, 100,000, 250,000 and 500,000 episodes. Some of the 19 rules contained in  $KB$ , are presented on Table 1.

```

(rule_1
  (if (logic (robot zone not_safe))
    (logic (target quadrant r1))
    (logic (obstacle quadrant r1)))
  (then (logic (robo action east)))(cf 0.84))
  ...
(rule_17
  (if (logic (robot zone safe))
    (logic (target direction ?x))
    (then (logic (robot action ?x)))(cf 0.90))
(rule_18
  (if (logic (robot zone victory))
    (then (logic (robot action success)))(cf 0.83))
(rule_19
  (if (logic (robot zone colision))
    (then (logic (robot action stop)))(cf 0.85))
  
```

Table 1: Rules Base of Mobile Robot Navigation

From rule 1 to rule 17, the rules chooses the basic behaviour according to the robot zone, the goal position, and the closest obstacle position. Rules 18 and 19, both require the reactive level to stop the robot and inform the cognitive level about the rules conclusion. Rule 18 identifies that the current goal was satisfied. In this situation, another goal should be chosen by the cognitive level.

Also, for rule 19, which characterizes a state where the mobile robot is surrounded by mobile obstacles, the cognitive level chooses another goal. This choice implies that another knowledge base must be used to scape from surrounded state, or wait for the collision zone state became false, it means wait for an mobile obstacles move out. The received reward for an action of the optimal policy  $\pi^*$  is normalized and stored in the rule as the certainty factor for that rule. A file in text format is generated which contains the rules base, or in other words the knowledge to execute the human like locomotion for the desired goal  $g_i$ . This file is downloaded to the omnidirectional mobile robot where the agent uses the stored knowledge to control the mobile robot locomotion.

## 4 Experimental Results

Once the symbolic knowledge base was embedded into CAA, two kinds of mobile robot locomotion experiments were done. First with the omnidirectional mobile robot AxeBot, and than with the NaO Humanoid robot. In both experiments, a simulated environment was used at first, and then with real mobile robot. The first experiment, the simulated one, the execution scenario was composed by: the robot, one target, and from 3 (three)




Figure 4: A success path execution with 8 obstacles

to 8 (eight) obstacles (static and dynamic). The positions of the environmental elements are randomly chosen, including the ball that is the target, or the goal. One of these scenarios, where 8 (eight) obstacles are used, is shown at figure 4.

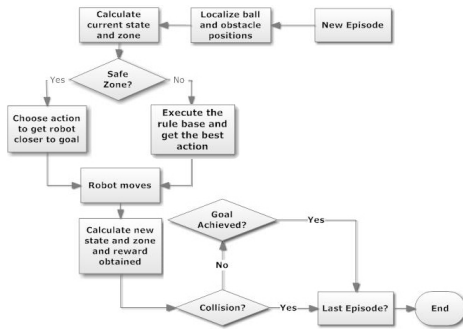


Figure 5: Testing Process

For each symbolic knowledge base,  $KB_i$ , 500 (five hundred) scenarios, which contain, from 3 (tree) to 9 (nine) obstacles, and the ball, randomly positioned, were tested. The effectiveness of each different policy automatically coded into the symbolic knowledge bases is presented at 6. The simulation algorithm is seen in figure 5.




Figure 6: Number of obstacles vs. Success rate

Clearly, the knowledge base generated from the 500,000 iterations presents, award majority, a better effectiveness than the other ones.

The second step seeks to validate the developed methodology for a mobile robot. The file which contains a knowledge base, generated by the 500,000 iteration learning process, was loaded into the CAA instinctive level embedded on the omnidirectional mobile robot. This file stores the mo-

tion plan, combining basic reactive behaviours, to move the robot to the current ball position, which contains 19 rules. Three different scenario, where used for this experiment. The respective trajectory tracking for one of these scenarios and the temporal evolution of robot positions can be seen at figure 7. The path was correctly planned and the real robot was able to track it under the real-time constraints, support the proposed methodology.




Figure 7: Temporal Evolution of Positions and Mobile Robot Trajectory

In the second experiment a new reactive level, which embedded a Zero Moment Point (ZMP) motion controller for humanoid robots was implemented. It was encapsulated in a process which communicates with the instinctive level using udp sockets. This new reactive process also encapsulates the local vision system, see Figure 8. The same exchange message pattern between the reactive level and the instinctive level was kept. The reactive level sends to instinctive level, the robot pose, the target pose and the position of the seen obstacles. The instinctive level identifies the current state and choose the behaviour that is executed by the ZMP motion controller.

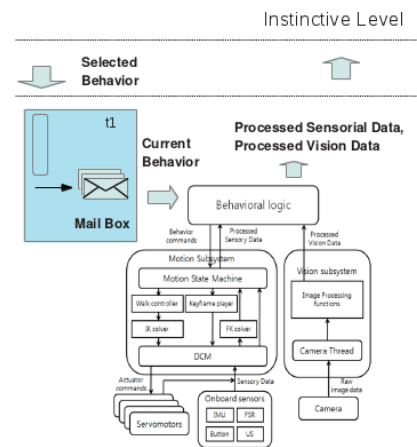


Figure 8: Temporal Evolution of Positions and Humanoid Robot Trajectory

The same Three different scenario, used at the first experiment with the omnidirectional mobile robot were successful repeated. The respective trajectory tracking for one of these scenarios and the temporal evolution of robot positions can be seen at Figure 9. Again, the path was correctly

planned and the real robot was able to track it.




Figure 9: Real Navigation using AxeBot

## 5 Conclusions and Future Works

A new methodology for knowledge-based agent to learn the knowledge from their interactions with the environment is presented in this paper. The new methodology combines Reinforcement Learning and Knowledge-Based Systems. The Reinforcement Learning obtains the optimal policy which is automatically coded using a knowledge representation formalism based on first order logic. The generated knowledge base is used by a cognitive agent embedded into the omnidirectional mobile robot that executes the learned human like locomotion task.

Considering only the closest obstacle, the state space of the environment is reduced drastically which accelerates the convergence to the optimal policy and mainly allows the proposed methodology to generate small knowledge bases appropriate to be embedded into mobile robots. The experimental tests realized show the planned points in real-time by the Instinctive level were received and executed successfully by the Reactive level.

For a different goal  $g_j$ , for example dribbling around an opponent, a new learning process would be necessary. This new learning process would generate a new knowledge base  $KB_j$ , which, once added to mobile robot, would make it able to achieve both  $g_i$  and  $g_j$ , using  $KB_i$  to achieve  $g_i$  to achieve  $g_j$ . In this sense, where a specific knowledge base is generated for a given task, it is possible see the proposed methodology as a symbolic knowledge learning methodology. Of course a good discussion about symbolic knowledge learning can be further researched in this situation, and we hope to support it with future additional results where the robot increases its skill with additional learned knowledge bases.

## References

Bittencourt, G. (2007). An embodied model for higher-level cognition, *29th Annual*

*Meeting of the Cognitive Science Society (CogSci2007)*, pp. 106–115.

Brooks, R. A. (1991). Intelligence without representation, *Artificial Intelligence* **47**: 139–159.

da Costa, A. L. and Bittencourt, G. (1999). From a concurrent architecture to a concurrent autonomous agents architecture, *Springer Lecture Notes in Artificial Intelligence 1856: 3rd International Workshop in RoboCup*, pp. 85–95.

Fatmi, A., Yahmadi, A. A., Khrijji, L. and Masmoudi, N. (2006). A fuzzy logic based navigation of a mobile robot, *World Academy of Science, Engineering and Technology* **22**: 169–174.

Jaradat, M., Al-Rousan, M. and Quadan, L. (2011). Reinforcement based mobile robots navigation in dynamic environments, *Robotics and Computer-Integrated Manufacturing* **27**: 135–149.

Manikas, T.W., A. K. W. R. (2007). Genetic algorithms for autonomous robot navigation, *IEEE Instrumentation and Measurement Magazine* **10**(6): 26–31.

Mccarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence, *Machine Intelligence* **4**: 142–152.

McGill, S., Brindza, J., Yi, S. and Lee, D. (2010). Unified humanoid robotic software platform, *5th Workshop on Humanoid Soccer Robots(Humanoids 2010)*, pp. 07–11.

Ouarda, H. (2010). A neural network based navigation for intelligent autonomous mobile robots, *International Journal of Mathematical Models and Methods in Applied Sciences* **4**(3): 177–186.

Pradhan, S., Parhi, D. and Panda, A. (2007). Navigation of multiple mobile robots using rule-based-neuro-fuzzy technique, *IEEE Instrumentation and Measurement Magazine* **34**(2): 142–152.

Shayestegan, M. and Marhaban, M. (2012). Mobile robot safe navigation in unknown environment, *IEEE International Conference on Control System, Computing and Engineering*, pp. 44–49.

Torrey, L. (2009). *Relational Transfer in Reinforcement Learning*, PhD thesis, University of Wisconsin-Madison, Computer Sciences Department.