A Cognitive Architecture for Instrumental Learning in Smart Agents

Wandemberg Gibaut[1], Ricardo R. Gudwin[1]
[1]Department of Computer Engineering and Industrial Automation (DCA), FEEC, UNICAMP

Introduction and Hypothesis:
The research field of Cognitive Architectures relates to getting inspiration from the different models of the human brain / human mind coming from neuroscience and neuropsychology and building smart artificial agents able to reproduce the cognitive capabilities found in animals. Among these capabilities, Instrumental Learning is a kind of learning where an agent learns from observing the results of its own actions on the environment.

Objective:
The objective of the present work was to propose a Cognitive Architecture, grounded on findings from neuroscience, for the smart control of an artificial creature in a computer game with a high degree of freedom on its actions. The chosen platform, Minecraft, available through the Mälmo Platform [3], is a 3D environment where, due to a huge state space, conventional control techniques are usually not suitable. In particular, we sought to analyze how models of Episodic Memory [2] may help in a cognitive agent's learning and decision making processes. To implement that, we used unsupervised learning techniques (Reinforcement Learning [1], for example) and Neural Networks.

Methods:
We took as methodology the exploration of the space state and checking the agent's learning. We used five types of controllers:

- Simple Look-up table Reinforcement Learning with ε-greedy policy: represents a classical method of decision-making process with guaranteed convergence.
- Cognitive controller purely based on Expectations: controller model whose decisions are simply based on a reinforcing structure with a neural network as approximator.
- Cognitive Controller with Episodic Memory: similar to the previous one, but using Memory to improve the decision-making process.
- Cognitive Controller with Exploration: the same as the previous one, but the agent decides to explore new situations in detriment of what it already knows.
- Complete Cognitive Controller: similar as previous, but this version can record sequences of actions and make plans. In each execution, the agent chooses to follow the best plan so far in 50% of cases and to explore new possibilities in the remaining ones.

To validate the full controller and the Architecture, and measure its performance, we did 10 rounds of 50 consecutive executions and defined some metrics to get an overview of the agent's learning with each controller. They are: the total number of victories per round, the average victory execution time in seconds and the average number of commands per victory.

Relevance:
This artificial mind project wraps together many known techniques in the research areas of Neurotechnology and Artificial Intelligence. This allows the Architecture to be used in a wide range of applications, including those which may benefit from autonomous learning, as drones, self-driving cars and brain-computer interfaces.
The proposed Architecture is also in accordance with other works present in the literature, aligning several theories in an unique framework, being potentially competitive with the existing ones.

References: [1] doi: 10.1017/S0269888999003082; [2] doi: 10.1037/h0020071; [3] JOHNSON, M. et al., Proceedings of the Twenty-Fifth IJCAI, 4246–424, 2016.