Planning, Exploration and Creativity in Motivational Agents Inhabiting Unknown, Dynamic Environments*

Luís Macedo[1,2]
[1]*Department of Informatics and Systems Engineering, Coimbra Engineering Institute, 3030 Coimbra, Portugal*
[2]*Artificial Intelligence Lab, Centre of Informatics and Systems, University of Coimbra, 3000 Coimbra, Portugal*
lmacedo@isec.pt; macedo@dei.uc.pt

Although, planning and exploration of unknown environments have been previously addressed in multi-agent environments (e.g.: [6]), we believe that in addition to these activities, agents may also benefit from exhibiting creativity [1] so that they are able to imagine or invent new things (objects) that may be helpful or simply pleasant for the agents that inhabit the environment. Psychological and neuroscience research over the past decades suggests that motivations (emotions, drives and other motivations) play a critical role in these activities that involve decision-making and action, by influencing a variety of cognitive processes (e.g., attention, perception, planning, etc.). Actually, on the one hand, recent research in neuroscience (e.g.: [2]) supports the importance of emotions on reasoning and decision-making. On the other hand, there are a few theories in psychology relating motivations (including drives and emotions) to action (e.g.: [4]).

In this thesis we study the role of motivations on planning, exploration and creativity in autonomous agents that populate unknown, dynamic environments. We have developed a multi-agent environment in which, in addition to inanimate agents (objects), there are two main kinds of animate agents interacting in a simple way: the creators, whose main function is to create things (objects, events), and the explorers whose goal is to explore the environment, analyzing, studding and evaluating it. In spite of this classification, there are agents that may exhibit the two activities: exploration and creation. In addition to these two activities, animate agents are able to generate plans. Planning plays a central role in reasoning/decision-making by supporting the other two activities (exploration and creativity). Actually, in our approach, creativity and exploration involve planning: when exploring the environment an agent has to plan a sequence of actions required to visit an unknown region or entity; when creating, an agent has to plan the sequence of actions required to come up with an original and valuable object.

The architecture of an agent includes the following modules: memory (for entities, plans, and maps of the environment), goals/intentions, desires, motivations, and reasoning/decision-making.

The memory of an agent stores information about the world. This information may be the configuration of the surrounding world such as the position of the entities (objects and other animated agents) that inhabit it, the description of these entities themselves, descriptions of the sequences of actions (plans) executed by those entities and resulting from their interaction, and, in generally, beliefs about the world. This information is stored in several memory components. Thus, there is a metric (grid-based) map to spatially model the surrounding physical environment of the agent. Descriptions of entities (physical structure and function) and plans are stored both in the episodic memory and in the semantic memory.

The module of motivations receives information from the current state of the environment and outputs the intensities of emotions, drives and other motivations (e.g.: surprise, curiosity, hunger, etc.). These feelings are of primary relevance to influence the behaviour of an agent.

Desires are states of the environment the agent would like to happen, i.e., they correspond to those states of the environment the agent prefers. This preference is implicitly represented in a mathematical function that evaluates states of the environment in

* The PhD of Luís Macedo is financially supported by PRODEP III.

terms of the positive and negative feelings they elicit in the agent. This function obeys to the Maximum Expected Utility principle [5].

Goals/intentions may be understood as something that an agent wants or has to do. These might be automatically generated by the agent or given by other agents. The reason for trying to achieve a goal might be because that achievement corresponds to a state of the environment that makes it feel positive feelings. Thus, goals may be thought of as a subset of the agent's desires.

The deliberative reasoning/decision-making module receives information from the internal/external world and outputs an action that has been selected for execution. The agent starts by computing the current world state. This is performed by generating expectations or assumptions for the gaps in the environment information provided by the sensors. Then, new intentions/goals (e.g.: *visitEntity*(y), *visitLoc*(x), *rechargeBatterry*(), *createObj*(z)) are generated and their Expected Utility (EU) computed based on the feelings that the agent may "feel". According to this EU, the set of goals of the agent are ranked, and the first one, i.e., the goal with highest EU is taken and a Hierarchical Task Network (HTN) plan (e.g.: [3]) is generated for it.

The planner is the core of the deliberative reasoning/decision-making module. The agent uses a planner that combines the technique of decision-theoretic planning with the methology of HTN planning in order to deal with uncertain, dynamic large-scale real-world domains. Unlike in regular HTN planning, the planner can generate plans in domains where there is no complete domain theory by using cases of previously successful plans instead of methods for task decomposition. It generates a variant of a HTN - a kind of AND/OR tree of probabilistic conditional tasks - that expresses all the possible ways to decompose an initial task network. The EU of alternative plans is computed beforehand at the time of building the HTN and it is based on the expected positive and negative feelings that the agent feels if the plan is executed. Plans that are expected to elicit more positive feelings (happiness, surprise, etc.) and less negative feelings (e.g.: hunger) are assigned a higher EU.

Unlike planning that is directly a part of the deliberative reasoning/decision-making module, exploration and creativity are activities that result from it, depending on the kind of goals generated.

When performing exploration, the aim of an agent is twofold: (i) acquisition of maps of the environment – metric maps – to be stored in memory and where the cells occupied by the entities that populate that environment are represented; (ii) construction of models of those entities. Exploration may be performed by single or multiple agents. Each agent autonomously generates goals for visiting unknown entities or regions of the environment (goals of kind *visitEntity* or *visitLoc*) and builds a HTN plan for each one. Goals and plans that are expected to cause more positive feelings and less negative feelings are preferred. Thus, each agent performs directed exploration using an action selection method based on the maximization of the intensity of positive feelings and minimization of negative ones. Relevant motivations for directing exploration are for instance curiosity, surprise and hunger. The exploration strategy for multiple agents relies on considering a team leader that, based on the information provided to it by the members of the team as they perform their single exploration, builds a joint metric map, a joint episodic memory and a joint plan in order to be shared by all the members of the team.

When performing creativity, an agent generates goals for the creation of novel, original and valuable entities (goals of kind *createObj*) and builds a plan for each one. Like in exploration, goals and plans that are expected to cause more positive feelings and less negative feelings are preferred. Motivations such as surprise and curiosity that capture variables such as novelty or unexpectedness, respectively, are hence important for creativity.

Several experiments have been performed to evaluate the role of motivations on planning, exploration and creativity.

Main References

[1] M. Boden, *The Creative Mind: Myths and Mechanisms*. New York: Basic Books, 1992.

[2] A. Damásio, *Descartes'error, Emotion Reason and the Human Brain*. New York: Grosset/Putnam Books, 1994.

[3] K. Erol, J. Hendler, and D. Nau, "UMCP: A sound and complete procedure for hierarchical task-network planning" in *Proceedings of the International Conference on AI Planning Systems*, 1994, pp. 249-254.

[4] C. Izard, *The Psychology of Emotions*. NY: Plenum Press, 1991.

[5] S. Russel and P. Norvig, *Artificial Intelligence-A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall, 1995.

[6] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes, "Coordination for Multi-Robot Exploration and Mapping" in *Proceedings of the AAAI-2000*, 2000.