ACHIEVING CREATIVE BEHAVIOR USING CURIOUS LEARNING AGENTS

Creativity is often associated with surprise, novelty, usefulness, and value. These characteristics do not assist in the development of a model for intelligent systems to achieve creative behavior, since they are characteristics that help identify when something or someone has been creative, as post-facto evaluation. Rather, models of creative behavior for intelligent systems draw on process models such as analogical reasoning and induction, or on principles such as "make the familiar strange" or "make the strange familiar". In this paper we describe a computational model of curiosity, based on cognitive models of novelty and interest, as a focus of attention in learning algorithms. We show how this combination of curiosity and learning can be the core reasoning processes in agent-based systems to achieve creative behaviors.

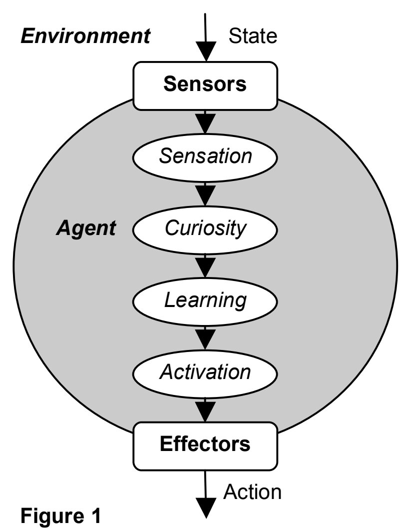

Figure 1

In our curious learning agent, the reasoning process is decomposed into four sub-processes: sensation, curiosity, learning and activation, as shown in Figure 1. The sensation process transforms raw data about the environment into structures to facilitate further reasoning – such as the observed state of the environment, the change or 'event' between the current and previous observed states, and actions performed by humans or other agents. The curiosity process reasons about the current observed state and events to identify an 'interesting' goal, which is used to focus the learning process. The learning process performs a learning update to incorporate the previously observed state, action and current rewards into a policy defining how the agent should act. Finally, the activation process selects an action to perform from the learned policy. This model can be used in both reinforcement and supervised learning scenarios.

The role of curiosity in reinforcement learning is to focus the learner's attention on an interesting goal long enough for a policy to be learned. The value computed by the curiosity process becomes the reward signal for the learning process. Our curiosity process comprises two functions: Stanley's model of habituation for computing novelty (Stanley, 1976); and the Wundt curve for computing interest (Wundt, 1910). We have developed our curious learning model in the context of generating creative behaviors for non-player characters in games. These characters are more adaptive and dynamic than characters using traditional, rule-based approaches.

While reinforcement learning can be an appropriate choice for generating creative behaviors in game environments, in other domains learning by trial and error is inappropriate. For example, one of the key goals of intelligent rooms research is to build environments – such as meeting rooms – that can unobtrusively adapt to the changing needs of their inhabitants. In this case, a supervised learning approach is more appropriate. The role of curiosity in supervised learning is to determine which examples of behaviors performed by humans should trigger learning and which observed states should trigger action. We are developing and trialing a curious supervised learner in the Virtual Sentient Room, an intelligent meeting room in *Second Life*, modeled on a real world university meeting room, to demonstrate the benefits of a curious room.

The simulations described in this paper show how curious learning agents are able to exhibit creative behaviors in different learning scenarios by focusing attention on "curious" events in the environment.