Learning to Interact:
Connecting Perception with Action in Virtual Environments
(Short Paper)

Pedro Sequeira
INESC-ID / DEI / IST
Av. Prof. Cavaco Silva
2780-990 Porto Salvo
Portugal
pedro.sequeira@gaips.inesc-id.pt

Ana Paiva
INESC-ID / DEI / IST
Av. Prof. Cavaco Silva
2780-990 Porto Salvo
Portugal
ana.paiva@inesc-id.pt

ABSTRACT
Modeling synthetic characters which interact with objects in dynamic virtual worlds is important when we want the agents to act in an autonomous and non-preplanned way. Such interactions with objects would allow the synthetic characters to behave in a more believable way. Once objects offer innumerous uses, it is essential that the agent is able to acquire the necessary knowledge to identify action possibilities in the objects while interacting with them. We propose a conceptual framework that allows the agents to identify possible interactions with objects based in past experiences with other objects. Starting from sensory patterns collected during interactions with objects, the agent is able to acquire conceptual knowledge about regularities of the world, its internal states and its own actions. The presented work also proposes that such acquired knowledge may be used by the agent in order to satisfy its needs and goals by interacting with objects. Preliminary tests were made and it is possible to state that our agents are able to acquire valid conceptual knowledge about the regularities in the environment and its objects, its own actions and causal relations between them.

Categories and Subject Descriptors
I.2.6 [**Artificial Intelligence**]: Learning – *concept learning, knowledge acquisition*

General Terms
Algorithms, Experimentation, Human Factors

Keywords
Learning agents, believable qualities, object interaction, affordances, perception

1. INTRODUCTION
In the past few years 3D virtual worlds have had an enormous growth mainly due to the success of computer games. These worlds are often populated by autonomous synthetic agents capable of acting and interacting with their environment, with other agents and also with humans. The virtual worlds are composed of several objects that may have different types, uses and effects in the world. In order for the characters to fully take advantage of the objects they perceive in the world, they should "know" them and their consequences beforehand. This is possible when the number of objects is limited and the characters are partially scripted to know all kinds of objects and their uses.

However, the support-mechanisms for a system that deals with agent-object interactions in virtual worlds grow to be very complex especially when the number of both objects and their available interactions becomes bigger and perhaps dynamic. Such system has to allow the characters not only to interact with their environment in a non-predictable way, but also to learn from their experiences and actions, thus creating unique personalities within the whole set of interacting agents. As such, the main questions we try to answer in the research here presented are:

How can an agent identify the possibilities of interaction with an object and the consequences of that interaction, based on previous past experiences with other objects?

How can these previous experiences guide the agent's behavior by making him choose the best interactions with objects in order to satisfy its current goals and needs?

In a previous work [7], we've presented a framework and implemented a test suite that allowed an agent to interact with its world and to learn concepts about objects and causal relations between action execution and changes on the agent's state. This framework allowed us to validate some learning methods to allow agents of acquiring conceptual knowledge.

In this paper we propose a conceptual framework including the agent's architecture in order to answer the questions before-mentioned. The solution is based in the qualitative and symbolic description of the objects' perceptual features perceived by the agent through time by interacting with its environment. The suggested solution allows pertinent actions performed by the agent to be memorized and action sequences to be created. We believe this will endow an agent with real-time planning capabilities and to be able to predict the usual results of its actions, as it is inspired in the notion of object affordances [5].

This paper is organized as follows. First we will present some related work which served as inspiration for the research here presented. Then we will describe the proposed framework to solve the presented problem, including the conceptual model and framework's architecture. Finally we draw some conclusions and provide the ideas of the future planned work.

Cite as: Learning to Interact: Connecting Perception with Action in Virtual Environments (Short Paper), Pedro Sequeira and Ana Paiva, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1257-1260.

2. RELATED WORK

Agent-object interactions is a research topic that can be seen at two different levels: one relating to several perspectives over object use and meaning acquisition; the other about current implemented systems that deal with agent-object interaction.

2.1 Perspectives over subject-object interactions

One way of looking at the interactions between agents and objects is to ground those interactions in the concept of affordances. The term affordance was first used by J. Gibson [5] to denote that each individual lives in a particular space of its environment, composed by the set of the affordances he acquires through time. Thus affordances are defined by the set of things that the environment offers an individual taking into account his physical capabilities, i.e., what the individual perceives.

Another important perspective is given by von Uexküll, who introduced the term functional tone [8]. For Uexküll, each individual ascribes a unique meaning to the objects it interacts daily, associating perceptions to actions through the creation of effector images. According to Uexküll the meanings that are invested upon the objects are influenced by the subject's current mental and emotional state and don't rely on their physical properties because they belong to each individual's perceptual world.

Differently, the main idea of the work developed by Zhang *et al.* [10] is that the generation of much of humans' intelligence results from processing information distributed across an individual's mind and external data derived from the world. The authors show that internal and external information about an individual can be interconnected to describe the environment's affordances, through the actions that are allowed to an individual to execute taking into account its intrinsic acting constraints. This way we can model different behaviors for different individuals.

One aspect that unites all these perspectives relates to the fact that the acquisition of action possibilities from objects relies in the subject's social and cultural background. Actions performed over objects make also part of the individual's experience by telling him what to expect when interacting with unknown objects that have perceptual features similar to previously sensed ones.

2.2 Systems dealing with agent-object interactions

In the context of our work and by looking at current systems that deal with agent-object interactions, we can divide them in systems that use plans to solve problems and systems that use affordances to guide the agents' behavior.

One way of dealing with the use of objects by agents is by attaching to the agent's minds the adequate knowledge to be used in planning for dealing with the problem of objects. In the area of action planning, Abaci *et al.* [1] propose the association of plans represented in a formal language to each available action in an object, thus providing the agent with semantic information about their consequences. In the SODAJACK system [4], Geib *et al.* present a planning module where the main idea is that a generic interaction with an object can generate several expansions according the agents' capacities and intentions, the characteristics of an object and the current state of the world. Each expansion forms a plan of actions that are transformed in motion directives in order for the agent to interact in the environment.

Some other systems have made efforts in trying to use the concept of Gibson's affordances to endow agents learning about facts in the environment by interacting with objects. This will allow them to satisfy its needs by performing specific actions. Within this area, one of the research works that greatly inspired the work here presented was done by Viezzer *et al.* [9] that proposed a way to internally represent affordances by joining external characteristics and possible uses for an object. This information can be used to select an object in the world that satisfies a certain agent's need. It is also proposed that a combination of internal and external sensor values describe the agent's state. A set of movement sensors can describe an action.

Another inspiring work is presented by Cohen *et al.* [2] which proposed a system that characterizes the objects by statistically analyzing sensory information perceived by the agent when interacting with them. This allows an agent to learn concepts of things and activities through interaction with its world. In another work [3], Cos and Hayes allow an agent to associate affordances with expectations of changing its internal state. This estimation can be refined as the agent interacts again with the same objects, allowing him to learn from past experiences.

Finally, another contribution for the work here presented comes from a recent seminar [6]. Some architectural issues were discussed in order to create an affordance-based system. According to the authors, a working implementation would enable a robot to find more action alternatives than pure appearance-based perception approaches. We believe that such advantages will also outcome from building synthetic characters inspired in affordance-based paradigms.

3. FRAMEWORK

The proposed solution is a framework inspired in the notion of affordances. We propose these possibilities of action to be perceived by the agent by sensing its environment and taking into account its current needs, goals, emotions, etc., i.e., its internal state. Through time, the agent is able to perceive some patterns in its sensory data and creates concepts relating facts that change or maintain often together. The proposed framework also considers the acquisition of the notion of causality by recording facts that happen before or after other facts thus creating relations between concepts. This conceptual knowledge and relations can later be used to guide the agent's behavior and help him fulfill its current needs and goals through the use of objects in its environment. The modeling of the agent's perceptions is based in Cohen's [2] work and the knowledge acquisition process and behavior control relies in Viezzer's [9] system and theories. Moreover, the proposed framework's architecture is based in some ideas taken from the Dagstuhl seminar [6].

3.1 Conceptual Model

The framework's main elements are ***Agents*** which explore a certain environment and ***Objects*** that are characterized by a set of perceptual features providing its entire sensory information.

Agents' sensors are described through a set of streams. A ***Stream*** is a sensory channel through which the agent receives internal or external stimuli, here called ***Tokens***. By interacting with objects, the perceived tokens will update the agent's streams. For example, if the agent sees an orange fruit, stream "color" will be updated containing the token "orange". A ***Sensation*** makes the connection between a certain stimulus and the sensory channel where it must be placed whenever the agent perceives something.

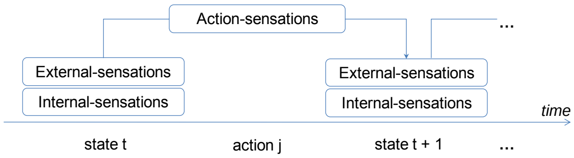

Figure 1 – Transition between agent's states by executing an action (adapted from [9]).

As we can see from Figure 1 there are 3 types of streams which give origin to 3 types of sensations. ***External-streams*** relate the agent to its environment and describe what it feels when sensing objects. ***Internal-streams*** relate the agent to itself, allowing it to perceive its current internal state. ***Action-streams*** are a special type of internal streams that allow the agent to perceive its own actions. They describe his conscience relating its own actions, like for example an agent sensing that his hand is closing.

An ***Action*** is the result of an internal process of the agent that can result in a variation of its internal or external state. Action execution provides information to the agent at two levels of perception: perceiving the action itself and perceiving the consequences of the action. Actions then define the agents' experience of activity. In Figure 1 we can see that at a certain instant of time, the agent's state is defined by its external and internal sensations. Transitions between states are described by action-sensations activated during the execution of certain action.

3.2 Architecture Overview

Our framework's architecture is a 3-layered agent architecture to manage agent-objects interactions in order to solve the presented problem, as we can see in Figure 2:

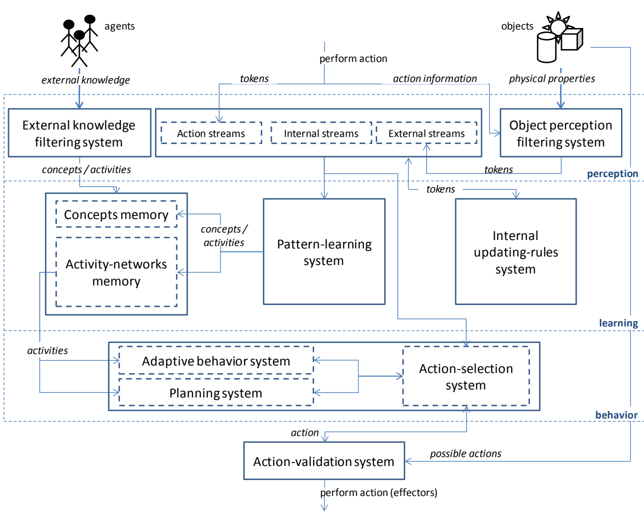

Figure 2 – Proposed framework's agent architecture.

The framework's objective is to ultimately produce conceptual knowledge and guide the agent's behavior by improving its acting performance. This is achieved by learning concepts relating the objects in the environment and the relations that the agent's actions have on those properties and on the agent's internal state. As such, the levels in the architecture are responsible to create, process and manipulate the information that will endow an agent of, staring from its perceptions, knowing what actions must it execute to fulfill its current goals.

At each iteration, the system executes the agent's pending actions by placing the correspondent action-tokens in the correct action-streams. Both ***external-knowledge*** and ***object perception systems*** check if new external knowledge has been acquired.

The ***internal updating-rules system*** updates certain internal-streams according to some values present in a certain internal variables. For example, we can imagine a rule that places the token "sleepy" when the agent is "awake" for some time.

As suggested by the framework's architecture, the agents' capabilities evolve according to two different phases: the first consists in gathering information from the environment and from the agent's usual activities; the second consists in shaping the agent's behavior by allowing him to perform actions that, according to the agent's mind, will lead him to "good" states.

3.2.1 Conceptual knowledge learning

The objective of this phase is to create and update activities stored in the agent's ***memory components***. The agent's learning evolution is achieved by monitoring the occurrence of events related with the presence of tokens in the agent's streams. If some events occur simultaneously and frequently (a significant number of times [2]), then some knowledge structures that relate them are created by the ***pattern-learning system***. As we can see from Figure 3, the agent starts by discovering concepts about the world, itself and its activities, and ends up building a set of regular activities in an activity-network, growing overtime both in size and complexity.

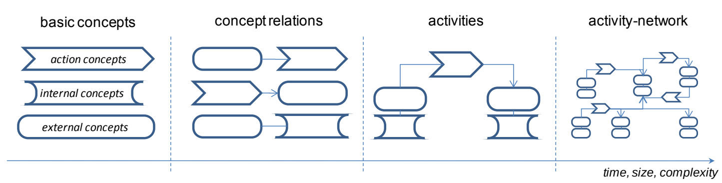

Figure 3 – Knowledge structures growth overtime.

The first kind of knowledge the agent learns is ***Concepts***, which are pieces of knowledge that the agent learns through time by interacting with its environment. It consists in a set of two or more sensations of the same type, i.e., a set of internal, external or action sensations. To see if such sensations are related (associated) we use statistical mechanisms [7] to determine if the set of sensations both start and stop a significant number of times, and also determine their association strength. This strength represents the agent's confidence in the existence of such concept in the world. As the system evolves the agent's confidence on the concepts will change, and some previously created concepts can even be "destroyed".

After the creation of concepts, relations between two types of concepts are generated. When the agent experiences certain action-sensations after perceiving some other sensations, the system can create the relation that tells that such agent's state potentiates performing that action. Similarly, relations between action-sensations and other internal or external-sensations tell us that the action usually enables the appearance of some agent's state.

These relations are the basis for the creation of larger structures of knowledge called activities. ***Activities*** represent patterns of actions comprising the agent's experience. It is created by the ***pattern-learning system*** and corresponds to the transformation from certain internal and/or external concepts (initial state) to other internal and/or external concepts (final state) by means of action concepts (action), as we can see in Figure 4:

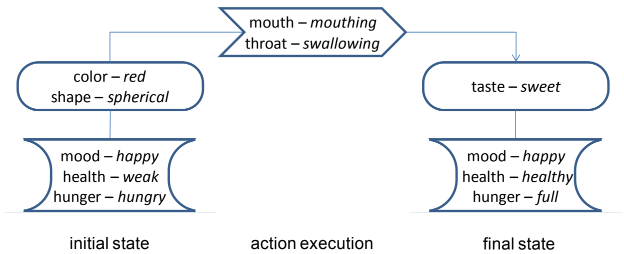

Figure 4 – Example of an activity acquired by the agent.

Besides the initial and final states and an action-concept, an activity has a certain probability associated. It represents the chance of, having the agent a current state matching the activity's initial state, by executing the action, ending up perceiving the sensations included in the final state.

As complex activities are gradually created, the system can begin to connect them with each other, storing activities in memory and associating them in ***Activity-networks***. It's important to notice that this link between actions that the agent performs in the environment and the learning of concepts about objects allows the agent to be able to experiment in the environment in a proactive manner, leading to goal-directed behavior once the relations are established.

3.2.2 Guiding the agent's behavior

Activity-networks are the elements stored in memory that allow the ***planning mechanism*** to create plans of action that the agent can follow to achieve some desired state. Activities can be seen as "IF *state n* AND *action j* THEN *state n+1*" rules. As such, the system is able to look within the agent's ***activity-networks memory*** for a chain of actions that, taking into account its current state, will lead with some probability to the desired state:

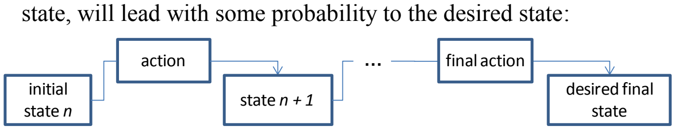

Figure 5 – Creation of plans of by chaining activities.

This kind of behavior allows the improvement of the agent's action performance and autonomy. Activity-networks combined with a planned behavior allow the agent to somehow look a few steps further into its future and guess what the results of performing some actions probably are.

4. CONCLUSIONS AND FUTURE WORK

We have presented a conceptual framework that proposes the connection between perceptions and the execution of actions as the basis to improve agent's acting performance and autonomy in virtual worlds. We believe that such capacity is possible by analyzing sensory data gathered from the environment while the agent interacts with it. Patterns discovered in this data through time allow the agent of learning conceptual knowledge about the world objects, its internal states, its actions, and causal relations between them. In this manner, the agent is able to adapt to new situations and worlds having into account its past experiences.

Like humans, synthetic characters don't inhabit virtual environments by themselves. They live in societies, exchanging information and creating relationships with other agents. Currently our framework is designed to directly receive external knowledge representing the agents' social and cultural beliefs. As each agent's knowledge is based in its own experiences, we believe that such exchange of information will allow us to build more believable agents as unique personalities emerge. In the future, we would also like to test the possibility of this cultural and social knowledge to be acquired by an agent by observing other agents in their activities just like it learns from its actions.

5. ACKNOWLEDGEMENTS

This paper is supported by the eCIRCUS (Contract no. IST-4-027656-STP) project carried out with the provision of the European Community in the Framework VI Programme. The authors are solely responsible for the content of this publication. It does not represent the opinion of the European Community which is not responsible for any use that might be made of data appearing therein.

6. REFERENCES

[1] Abaci, T., Ciger, J. and Thalmann, D. 2005. Planning with Smart Objects. *Int. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG 2005)*

[2] Cohen, P. R., Atkin, M., Oates, T. and Beal, C. R. 1997. Neo: Learning conceptual knowledge by sensorimotor interaction with an environment. In *Proceedings of the First International Conference on Autonomous Agents*.

[3] Cos, I. and Hayes, G. 2002. Behaviour control using a functional and emotional model. In *Proceedings of the 7th Int. Conf. on Simulation of Adaptive Behavior on From Animals To Animats* (Aug. 05 - 09, 2002). MIT Press.

[4] Geib C. W., Levison L. and Moore M. B. 1994. *SodaJack: an architecture for agents that search for and manipulate objects*. Technical Report. Department of Computer and Information Science, University of Pennsylvania.

[5] Gibson, James J. 1979 *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA.

[6] Rome, E., Hertzberg, J., Dorffner, G., Doherty, P. et al. 2006 Towards Affordance-based Robot Control - Abstracts Collection of the Dagstuhl Seminar 06231.

[7] Sequeira, P., Vala, M. and Paiva, A. 2007. "What can I do with this?" – Finding Possible Interactions between Characters and Objects. In *Proc. of the 6th Int. Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)* (Honolulu, USA, May 14-18, 2007).

[8] von Uexküll, J. 1928. *Theoretische Biologie*. Berlin. Springer Verlag.

[9] Viezzer, M. and Nieywenhuis, C. H. M. 2005. Learning affordance concepts: Some seminal ideas. *International Joint Conference on Artificial Intelligence*, Edinburgh.

[10] Zhang, J. and Norman, D. A. 1994. Representations in Distributed Cognitive Tasks. In *Cognitive Science*, 18(1).