"What can I do with this?"

Finding possible interactions between characters and objects

Pedro Sequeira
INESC-ID / DEI / IST
Av. Prof. Cavaco Silva
2780-990 Porto Salvo
Portugal

pedro.sequeira@gaips.inesc-id.pt

Marco Vala
INESC-ID / DEI / IST
Av. Prof. Cavaco Silva
2780-990 Porto Salvo
Portugal

marco.vala@inesc-id.pt

Ana Paiva
INESC-ID / DEI / IST
Av. Prof. Cavaco Silva
2780-990 Porto Salvo
Portugal

ana.paiva@inesc-id.pt

ABSTRACT
Virtual environments are often populated by autonomous synthetic agents capable of acting and interacting with other agents as well as with humans. These virtual worlds also include objects that may have different uses and types of interactions. As such, these agents need to identify possible interactions with the objects in the environment and measure the consequences of these interactions. This is particularly difficult when the agents never interacted with some of the objects beforehand. This paper describes *SOTAI – Smart ObjecT-Agent Interaction*, a framework that will help agents to identify possible interactions with unknown objects based on their past experiences. In SOTAI, agents can learn world regularities, like object attributes and frequent relations between attributes. They gather qualitative symbolic descriptions from their sensorial data when interacting with objects and perform inductive reasoning to acquire concepts about them. We implemented an initial case study and the results show that our agents are able to acquire valid conceptual knowledge.

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

General Terms
Algorithms, Experimentation, Human Factors

Keywords
Synthetic agents: human-like, lifelike, and believable qualities, learning agents

1. INTRODUCTION
Virtual environments are often populated by autonomous synthetic agents capable of acting and interacting with other agents as well as with humans. Such virtual environments also include objects that may have different types, uses and effects in

the world. As such, in order for the characters to fully take advantage of the objects they perceive in the world, they should "know" them beforehand. Of course 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 in the world. However, as the complexity of the virtual worlds increase and the types of objects become large and perhaps dynamic, other approaches need to be followed. On possible way would be for the objects to embed information about themselves that they pass to the agents as they interact with them. Such an approach (Smart-Objects) is indeed quite flexible and allows for a decentralization of the knowledge about objects detaching it from the agent's minds. However, the process of creating such objects is difficult, and as the characters become more believable and autonomous, part of that knowledge should be acquired by the agents.

So, the main question underlying the research here is presented is: "*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?*"

Inspired by the way we humans learn how to use objects, and by the notion of *affordances* by Gibson [11] we have developed a framework that will help agents to identify possible interactions with unknown objects based on their past experiences. The objects in the world exist, per se, and it is by experimenting with them that characters learn their properties and extract the right uses for them. As such, agents can learn world regularities, like object attributes and frequent relations between attributes. They gather qualitative symbolic descriptions from their sensorial data when interacting with objects and perform inductive reasoning to acquire concepts about them.

This paper is organized as follows. First we will describe some related work that serve as inspiration for our research. Then we will describe the conceptual model underlying the SOTAI framework, describing its main features and how it was implemented. We continue by describing a small test case and show some of the results attained with it. Finally we draw some conclusions and provide some idea of the future planned work.

2. RELATED WORK
The area of object-agent interaction can be seen through four different perspectives. In this section we will try to cover these different approaches and summarize their influence on our approach.

2.1 Objects and world modeling

Kallman *et al.* [14][15] introduced the widely used paradigm of Smart-Objects. Their aim was to solve problems related to the graphical aspect of interactions between synthetic characters and objects. The main idea is to encapsulate within the object descriptions of its characteristics, properties, behaviors and all the necessary scripts associated with each possible interaction with it. As such, aspects like animation, including all the movements a character has to execute during an interaction are attached to the smart-object. The main advantages of this approach are the decentralization of the simulation control from the main animation control into the object itself, and the reuse of the objects between applications by employing a description that is independent of the application.

Also concerned with the modeling of the objects, Forbus [10] adopts a qualitative process theory, stating that and phenomena that occur in the world are a result of physical processes. This approach was recently applied in virtual world's simulation by Cavazza *et al.* [3], more precisely to model objects' behavior. A framework supporting qualitative processes was implemented, where objects were modeled using qualitative variables and physical states. The transitions between states were represented graphically by a certain animation.

2.2 Planning

Another way of dealing with the use of objects by agents is by attaching to the agent's minds the adequate knowledge and use planning for dealing with the problem of objects. This approach can be addressed in two ways in two categories: action planning and using pre-defined plans.

In the area of action planning, the work of Abaci *et al.* [1][2] aimed at solving a limitation found in Kallman's Smart-Object implementation which come from the fact that the available interaction plans were fixed. Such fact reduces the agent's ability of adapting to new situations. As a result, an extension was proposed to the Smart-Objects architecture, where to each possible action for an object has associated plans represented in a formal language. These plans provide the agents with semantic information about the respective actions through the description of the actions' consequences in the world as well as the agent's states.

In the SODAJACK system [16][17], Geib *et al.* present a planning module that makes the bridge between action goals that must be achieved by agents and the movements they have to perform when considering an interaction with a specific object, transforming task-actions in action directives. The motivation for this work comes from the idea that a generic task can generate several expansions according the agents' capacities and intentions, the characteristics of an object and the current state of the world.

2.3 Object categorization

Another aspect that has been addressed in the quest for a good model of Agent-Object interaction is the categorization of the objects. Gonçalves *et al.* [12][13] show how object categorization can be done in virtual worlds, by encoding simulated sensorial information taken from the object. Using this information as input to neural networks, they show that it is possible for an agent to index objects in its environment and interact with them. This system makes the connection between individual perceptions and actions.

Cos-Aguilera *et al.* [8][9] demonstrate a method that allows an agent of "learning that an object exhibiting certain regularities offers the possibility of performing a particular action". Such learning is done through interaction episodes with each object. The idea is to use the perceptual capabilities of the agent to classify regularities in its sensory space (grouping similar objects in clusters by means of a growing neural gas network). This allows the match between sets of regularities and potentials of action.

2.4 Affordances (possible uses)

Affordance has been defined by Gibson [11] as: "Offerings or action possibilities in the environment in relation to the action capabilities of an actor. Affordances are independent of the actor's experience, knowledge, culture, or ability to perceive. Their Existence is binary – an affordance exists or it does not exist". This definition of affordance well known in design communities has also been used in the areas of conceptual knowledge acquisition and using affordances to guide action.

One of the big inspirations of the work here presented was done by Cohen *et al.* [6][7] who tries to demonstrate how an agent can learn concepts through the interaction with a simulated virtual environment. By beginning with little initial structured information, it is possible for an agent to learn representations about objects, activities, places and other aspects relating the agent and its environment. It is discussed that object categorization is based not in objective features such as size, color or shape, but rather in interactive properties such as "graspable". These kinds of properties indicate the way the agent interacts with its environment.

Viezzer *et al.* [19][20] explore an architecture for the acquisition of affordance concepts. It is stated that when the animal-environment system is observed in its totality, we notice that animals have evolved in a way that allows them to detect properties of the environment which are relevant to their survival. Representations in memories are patterns of action derived on environment properties, in conjunction with patterns of interactions based on past experiences. In this sense, for example the concept of food can be acquired through a system that relates the action of eating with the expectation of an internal energetic increase.

As it will become clear along the next section, SOTAI was inspired by this previous research. The system allows for artificial agents to learn concepts about objects and events which occur in its environment. A conceptual model is presented and an analysis is made concerning related work.

3. SOTAI: A FRAMEWORK FOR MANAGING OBJECT-AGENT INTERACTIONS

Imagine Bob, an autonomous agent that is placed into a virtual environment about which he knows nothing. For Bob, the objects that he perceives are unknown as there haven't been any past experiences with them. Bob is hungry and as such he needs to eat to survive. The objects in the world are shown in Figure 1.
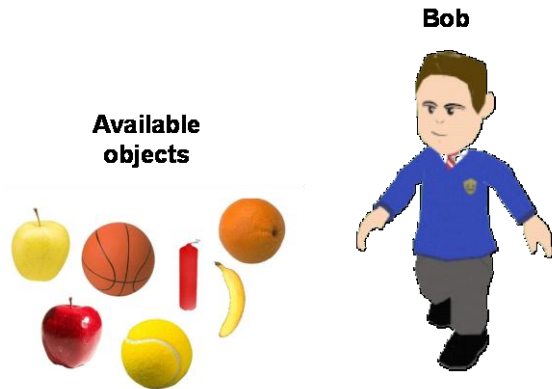
Figure 1 – "What object can Bob choose to eat?"

The framework here described, SOTAI (Smart ObjecT-Agent Interaction) is based on the idea that each interaction between agents and objects should provide some information about the facts that have changed in the world and in the agent caused by that interaction.

The framework is based on Gibson's idea of affordance presented in the early section. Its final goal is to allow artificial agents to recognize the possible interactions with an object, i.e. what it affords, taking into account past interactions with different objects. What our current model allows is the perception of world regularities by agents: facts that change or maintain often together and that can be grouped in concepts. The framework also considers the acquisition of the notion of causality by recording facts that happen before or after other facts and how frequently.

The acquisition and way of modeling world objects is based in Cohen's [6] system, which makes that some of the terms and concepts here presented have their origin in Cohen's work.

3.1 SOTAI Conceptual model

The SOTAI framework conceptual model can be divided into basic concepts and concepts that represent the conceptual knowledge an agent learns during the interaction with its environment.

3.1.1 Basic concepts

The main concept in our framework is the **Environment**, which is a place where *Agents* and *Objects* co-exist. An **Object** is a world entity with which an *Agent* can interact, and it is characterized by a group of *Actions* encoding the entire object's sensorial information. An **Agent** is an entity that has a group of sensory channels (here named *Streams*) and the ability to perform actions which change its internal *Streams*.

A **Stream** is an agent's sensor by which it receives internal or external stimulus. *Streams* can be seen as bags that at some moment can have one or more stimuli (the *Tokens*). For example, one *Stream* can be the smell sense. A *Stream changes* at a certain moment of time when some of the stimulus that the stream contained in its bag in a previous considered instant is not present in the current one. Otherwise it *maintained* its current values.

A **Token** can be seen as a symbol representing internal or external stimuli, whether they come from internal *Streams* respecting the *Agent* or external *Streams* respecting the *Environment*. One *Token* can be for example a "vanilla smell".

The correspondence between a *Token* and its respective *Stream* is called a **Sensation**. These pairs tell in which *Stream* should a *Token* be placed when an *Action* is performed. A *Sensation* is said to be *active* in some moment of time when the stimuli it represents is placed in the respective *Agent*'s sensory channel (*Stream*), and *inactive* otherwise. We say that a *Sensation starts* when it comes *active* and ends when it becomes *inactive* in some moment.

An **Action** is the result of an internal process that can result in the change of the internal or external state. It's characterized by a group of *Sensations* that are transferred to the agent when the *Action* is performed.

Figure 2 represents a simple *Environment* with our agent called Bob with three streams (color, shape and smell) and an orange *Object* with two possible actions (see and smell).
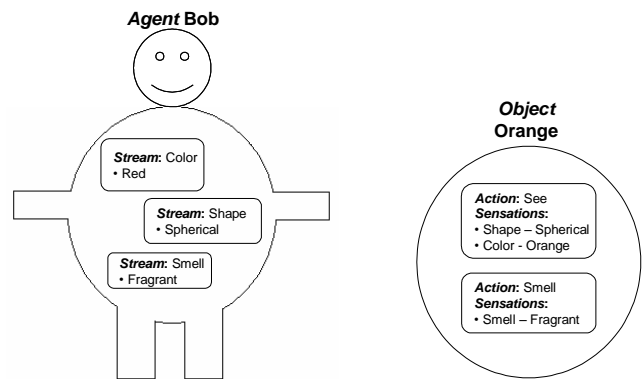

Figure 2 – an example of an *Environment* in SOTAI

3.1.2 Conceptual knowledge terms

In Figure 3 we can see the conceptual knowledge structures an agent apprehends by interacting with its environment. The smallest and first obtained structures are called *Scopes* and the largest and last acquired structures are called *Chains* of concepts.
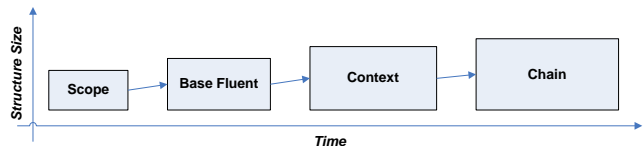

Figure 3 – conceptual structures and their relative sizes created for an agent by the SOTAI framework overtime.

The SOTAI system is updated in time steps. In each time step the system verifies if new knowledge structures can be created based on the current state of the system. To best describe how the system works and evolves and how the concepts are created by the SOTAI system overtime, we present a series of examples of how the structures can be created.

Assuming that the initial *Environment* state in *step=0* is the one illustrated in Figure 2, lets now assume that in *step=1 Agent* Bob performs *Action see* over the *Object orange*. The *Environment* state at this time step is illustrated in Figure 4.
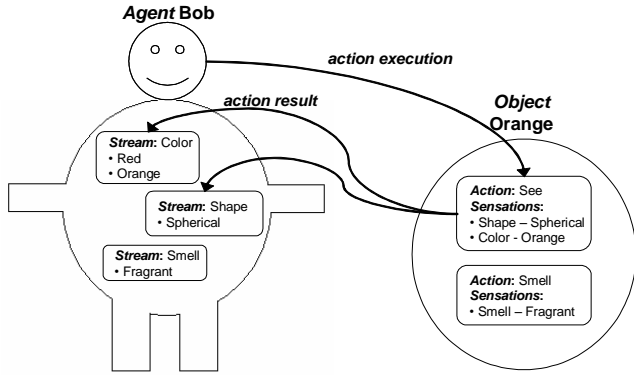
Figure 4 – the performing of *Action see* over *Object orange* and the result of that action on *Agent* Bob

The SOTAI system will initially create contingency tables between all pairs of streams to see what agent's streams change together often. In our example scenario, three contingency tables are then created between all possible pairs of the streams *color*, *shape* and *smell*. In *step=1* the tables would look like this:

	color *changed*	color *mantained*
shape *changed*	0	0
shape *mantained*	1	0

	smell *changed*	smell *mantained*
shape *changed*	0	0
shape *mantained*	0	1

	smell *changed*	smell *mantained*
color *changed*	0	1
color *mantained*	0	0

To determine if the *Streams* tested in the contingency tables are related, meaning they change or maintain together often, we first use a chi-squared statistical test using the formula:

$$\text{chi-squared} = \frac{N(O_{11} - E_{11})^2}{E_{11} E_{22}}$$

This formula only tells us whether the *Streams* are related. To measure their association strength we then use Phi-coefficient association measure, which is given using the formula:

$$Phi = \sqrt{\frac{X^2}{n}}$$

If this value exceeds a predetermined threshold, then we can say that the *Streams* are related and change or maintain together often, more than would be expected by chance if they were independent, as stated by Cohen *et al.*[6].

Continuing with our example, let's image that at a certain time step the middle table looks like this:

	smell *changed*	smell *mantained*
shape *changed*	13	18
shape *mantained*	10	504

Applying the statistical tests given earlier we can state that the streams *shape* and *smell* maintain together a significant number of times. So, we can say that they are related, and form the first kind of conceptual structure, a "*Scope*". ***Scopes*** are pairs of *Streams* in which we can look for greater forms of relation. In our example, we can form the *Scope smell – shape*.

The SOTAI system then creates two contingency tables between all pairs of tokens within each *Stream* of the *Scope*. Following our example, two tables are created like the next ones. The first table checks whether the *Tokens spherical* and *fragrant start* together often. The second one tests if the *Tokens end* together frequently.

	spherical *started*	non-spherical *started*
fragrant *started*	18	4
non-fragrant *started*	9	400

	spherical *ended*	non-spherical *ended*
fragrant *ended*	45	10
non-fragrant *ended*	6	394

Using the same statistical tests as before, the system creates another conceptual structure called *Base Fluent* when it discovers that the two *Tokens* both *start* and *end* together often. ***Base Fluents*** are pairs of *Sensations* indicating that two stimuli are frequently making their presence in the *Agent*'s sensory channels. We say that a *Base Fluent* is *active* in some moment when both the *Sensations* are *active* in the same instant of time. We say that a *Base Fluent starts* when it becomes *active*, and that it *ends* when it becomes *inactive*. We also say that a *Base Fluent F1 starts in the context of* another *Base Fluent F2* whenever *F1 starts* and *F2* has been *active* for a certain period of time.

In our previous example, we can create the *Base Fluent F1* represented by:

F1 — Shape: spherical, Smell: fragrant

Let's now assume that in a certain time step, another *Base Fluent* named *F2* is created:

F2 — Health: healthy, Taste: sweet

SOTAI system creates two contingency tables like the following ones, between all pairs of *Base Fluents* of the *Agent*.

	F1 *starts*	F1 *doesn't start*
F2 *active*	18	300
F2 *inactive*	50	56

|	F2 *starts*	F2 *doesn't start*
F1 *active*	18	4
F1 *inactive*	9	400

The first table checks whether *F1 starts in the context of F2* and the second one if *F2 starts in the context of F1*.

When the system discovers that one *Base Fluent* starts in the context of another a significant number of times, it creates a *Context*. ***Contexts*** represent cause-effect relations between *Base Fluents*. It's created in the sense that if a *Base Fluent* starts several times when another *Base Fluent* is active, then the former *Base Fluent* must be the effect caused by the latter one.

Continuing with our example, we can create the *Context* represented as follows:
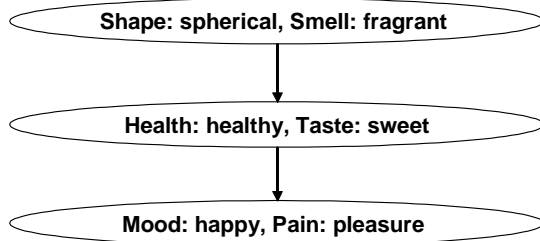
Context C1

F1 — Shape: spherical, Smell: fragrant

F2 — Health: healthy, Taste: sweet
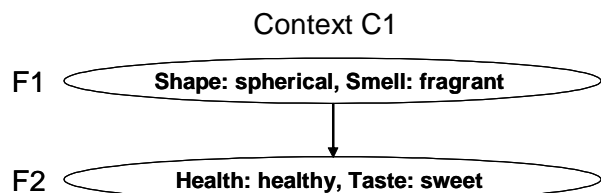
Context C1 tells us that *Base Fluent F1* may be causing the appearance of *Base Fluent F2*. Assume now that at a certain time step the following *Context C2* was created:
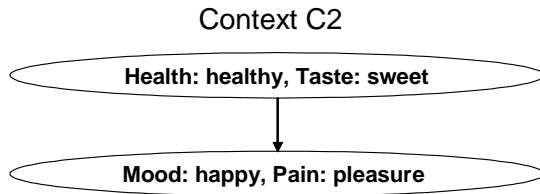
Context C2

When this happens, SOTAI system creates a structure called *Chain*. A **Chain** is a set of consecutive *Contexts* linked together and constitutes the largest knowledge structure an *Agent* can form. We can see them as paths in a cyclic graph where the nodes are *Base Fluents* and the arcs linking them are the *Contexts* between them.

Following the example, a *Chain* can be created looking like this:

This *Chain* can be seen as the following chain of events: when the agent sees something spherical and fragrant, its health improves and it eats something sweet. Later it becomes happy and pleasured.

Contexts linked together form graphs like the one described earlier. These graphs are called **Chain Lists**. Here's an example of a *Chain List* relating the example we've been presenting:
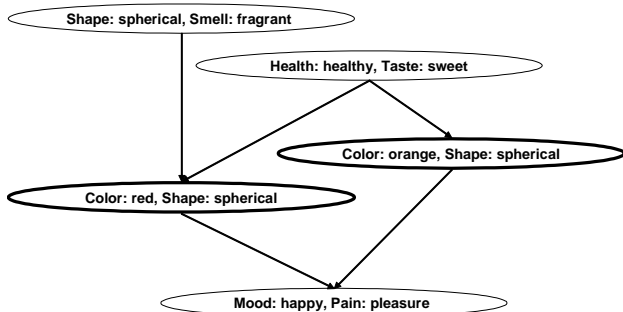
As we can see, the *Base Fluents* in bold both are caused by the same effects and both cause the same effects. When this happens, SOTAI groups the two *Base Fluents* into one structure called *Fluent*. A **Fluent** can then be defined as a set of *Sensations* that are related in the sense that they both appear as a result of the same events, and they both cause the same consequences in the *Agent* or the *Environment*. The operating of joining two *Base Fluents* into a larger *Fluent* is called **grouping**. *Grouping* allows the creation of larger conceptual structures from smaller ones. This operation is usually performed in a time stage where we think the *Agent* has learned most of its *Environment* regularities.

Finishing our example, the *grouping* operation would then result in the following *Chain List*:
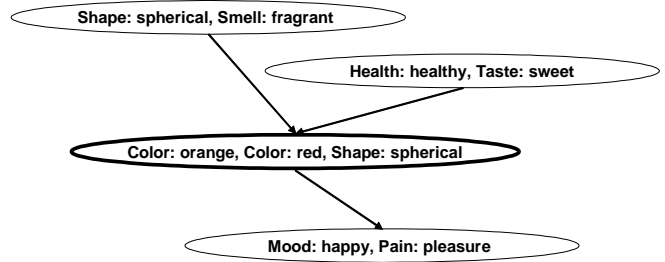
The next section describes a case study relating the SOTAI framework.

4. A CASE STUDY

The SOTAI framework conceptual model presented in the previous section was implemented in a small case study and we developed a test suite to gather results from the system.

4.1 Implementation

SOTAI system was implemented using *.NET Framework v1.1* and *C#*. The result was a dynamic library which implements the SOTAI system. Figure 5 summarizes the package diagram of the SOTAI library.
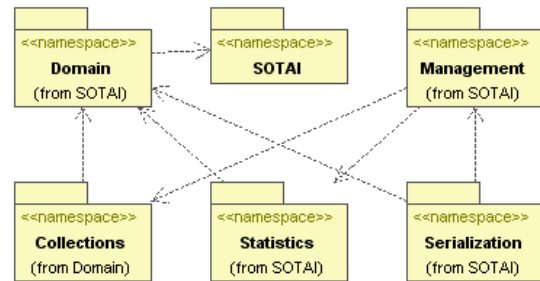

Figure 5 – UML diagram of SOTAI system's packages

The main packages are the Domain package which implements all the concepts from the conceptual model and the Management package which manages the conceptual structures creation and the contingency tables kept by the system.

4.2 Tests

We developed a test suite called *SOTAITester* which models an autonomous agent and a virtual environment containing 8 objects according to the SOTAI standard definitions. The objects positions in the environment are randomly generated at the beginning of each test.

The implemented agent has a reactive behavior. It explores his world and interacts with the objects. At each time step the agent smells, listens and looks at the objects according to its current position and distance to those objects. When close enough to an object, he touches and tastes that object, unless it's too big or causes too much pain to be eaten. The agent's object selection behavior is completely random and objective less. Initially the agent doesn't have any conceptual information about his environment. The agent has internal streams such as hunger, mood or health. External streams include sound, color, smell and pain among others. The application's user interface is illustrated in Figure 6.
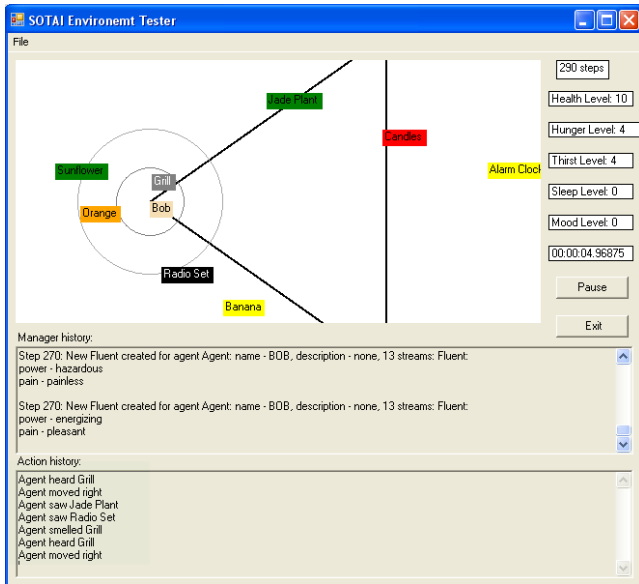
Figure 6 – *SOTAITester*'s user interface.

The application runs 30000 simulation steps. At each step the SOTAI system is updated and current actions (such as see, touch or mouth) are performed.

4.3 Results

A series of tests were made using the *SOTAITester* application according to the conditions described earlier. The main results consist of the largest and most frequent conceptual structures acquired through the simulation by the agent. Despite the random selection behavior, the agent usually learns most of its world regularities in less than 7000 simulation steps.

The largest *Fluents* learned by the agent were:

> **Hunger: hungry, Sound: frying, Smell: burnt**

that can be interpreted as the object *grill*, which makes the agent hungry,

> **Taste: metallic, Health: ill, Taste: bad, Pain: painful, Power: hazardous**

which can be interpreted as the things that are hazardous for the agent and that he doesn't like,

> **Taste: sweet, Health: healthy, Taste: pleasant, Pain: pleasant, Power: energizing**

which is the opposite of the previous concept. It represents the things that make the agent pleasured and healthy. This concept can be seen as a concept of food for the agent, something which makes him healthier;

> **Health: healthy, Hunger: full, Power: energizing**

This *Fluent* can represent the concept of an internal state of satisfaction for the agent.

5. CONCLUSIONS AND FUTURE WORK

We presented a framework to help agents identify possible interactions with unknown objects based on past experiences. Our approach is based on the notion of affordances which, from the point of view of our work, are interaction opportunities transmitted by objects to agents given their interaction possibilities.

So far we have focused on the acquisition of conceptual knowledge to detect similarities and cause-effect relations involving objects. We use sensorial streams to gather relevant information from the objects and then perform inductive reasoning to acquire concepts.

Our contribution enhances previous Cohen's [6] work because it features simultaneous acquisition of multiple symbols (*Tokens*) in the same sensorial channel and it only groups concepts (*Base Fluents*) into larger concepts (*Fluents*) when the agent has learned most of the environment's regularities (which is more efficient from a performance point of view). We also organize the information using the smart-object paradigm meaning that most of the interaction information is kept in the objects and the agent has only to know how to extract relevant information from them, which is done by experimentation.

In the future we plan to use concepts acquired from the objects to transmit interaction possibilities to the agents which will use this information to learn what to do with similar but unknown objects. A possible solution is to use reinforced learning algorithm, based on expectation and confidence, to guide the agent's behavior. The system will allow the association of levels of expectation between objects and certain actions. If an action performed with an object succeeds in a way that it satisfies the agent's expectations towards the pair object-action, the confidence level on that specific pair will rise. Otherwise if the action does not succeed, the confidence level will fall. Later, when the agent needs to satisfy certain necessities which involve a particular action, he can look for the object with the greatest confidence level to perform that action.

Having as an objective the integration of the SOTAI system with 3D virtual environments, tests will be made in order to find out if this integration is possible in a performance point of view and whether the realism needed for this kind of applications is not harmed with such integration. For the modeling of objects and the way the world evolves we can use the ideas of qualitative physics, whose integration in 3D environments has been widely used. This fact may give dynamic to the world, allowing objects and agents to evolve and have different states.

The importance of modeling the interactions between agents and objects in virtual worlds comes from the fact that such interactions show how the world changes in consequence of agents actions. As such, interactions are very important in the agent's learning process. They allow him to adapt to his world, observe the facts occurred and act in conformity. Through a continuous experimentation process with his world, an agent will be able to understand it and predict its behavior. SOTAI was born to model this behavior. It allows an agent to evolve through experimentation, discovering, acting, learning and living. We believe that this approach will allow synthetic characters in virtual worlds to become more and more autonomous and thus, more believable.

6. REFERENCES

[1] Abaci, T., Ciger, J., Thalmann, D.: "Action Semantics in Smart Objects" In: *Proceedings of the Workshop towards Semantic Virtual Environments*, Villars, Switzerland pp.121-126 (2005)

[2] Abaci, T., Ciger, J., Thalmann, D.: "Planning with Smart Objects" In: *WSCG (Short Papers)* 2005: 25-28

[3] Cavazza, M., Hartley, S., Lugrin, J., and Le Bras, M. 2004. "Qualitative physics in virtual environments". In *Proceedings of the 9th international Conference on intelligent User interface* (Funchal, Madeira, Portugal, January 13 - 16, 2004). IUI '04. ACM Press, New York, NY, 54-61. DOI= http://doi.acm.org/10.1145/964442.964454

[4] Cavazza, M., Charles F., Mead S., *"Agents' Interaction in Virtual Storytelling"*, Lecture Notes in Computer Science, Volume 2190, Jan 2001, Page 156

[5] Cavazza, M., Charles, F., Mead, S.J., "AI-based Animation for Interactive Storytelling", *Proceedings of Computer Animation*, 2001. IEEE Computer Society Press.

[6] Cohen, P. R., Atkin, M., Oates, T., and Beal, C. R. Neo: Learning conceptual knowledge by sensorimotor interaction with an environment. In *Proceedings of the First International Conference on Autonomous Agents* (1997), pp. 170--177.

[7] Cohen, P. R., Adams, N., Oates, T., and Beal, C. R. "Contentful mental states for robot baby". In *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, 2002.

[8] Cos-Aguilera, I., Cañamero, L., Hayes, G. M., "Learning Object Functionalities in the *Context* of Behaviour Selection", in *Proceedings of the 4th British Conference Towards Intelligent Autonomous Robots* (TIMR'03), August 2003

[9] Cos-Aguilera, I., Cañamero, L., Hayes, G. M., "Using a SOFM to learn Object Affordances", In *Proceedings of the 5th Workshop of Physical Agents* (WAF'04), Girona, Catalonia, Spain. March, 2004. DRAFT

[10] Forbus, K., "*Qualitative reasoning*". CRC Hand-book of Computer Science and Engineering. 1996, CRC Press.

[11] Gibson, J.J. (1979). "*The Ecological Approach to Visual Perception*". Boston: Houghton Mifflin.

[12] Gonçalves, L. M., Kallmann, M., and Thalmann, D., "Programming Behaviors with Local Perception and Smart Objects: An Approach to Solve Autonomous *Agents* Tasks" In *Proceedings of the XIV Brazilian Symposium on Computer Graphics and Image Processing* (Sibgrapi'01) (October 15 - 18, 2001). SIBGRAPI. IEEE Computer Society, Washington, DC, 184.

[13] Gonçalves, L. M., Kallmann, M., and Thalmann, D., "*Defining Behaviors for Autonomous Agents based on Local Perception and Smart Objects*" In Computer & Graphics, 26(6):887 897, December 2002

[14] Kallmann, M., Thalmann, D., "Modeling Objects for Interactive Tasks" In: *EGCAS'98 - 9th Eurographics Workshop on Animation and Simulation*, Lisbon (1998)

[15] Kallmann, M., "*Object Interaction in Real-Time Virtual Environments*", DSc Thesis number 2347, Swiss Federal Institute of Technology (EPFL), January 2001.

[16] Levison L., Geib C., Moore M., "*Sodajack: an architecture for agents that search for and manipulate objects*" In: Technical Report MS-CIS-94-16, University of Pennsylvania, Dept. of Computer and Information Science, Philadelphia, PA, 1994.

[17] Levison, L., "*Connecting Planning and Acting: Towards an Architecture for Object-Specific Reasoning*" In: Dissertation proposal, 1994.

[18] McGrenere J., Ho, W., "Affordances: Clarifying and evolving a concept". In *Proceedings of graphics interface* (2000, May) (pp. 179-186), Montreal, Canada, Retrieved September 1, 2003, from http://www.graphicsinterface.org/proceedings/2000/177/

[19] Viezzer, M. "*The Autonomous Acquisition of Affordance Concepts: Exploring a Neighbourhood in Design Space*", Technical Report CSRP-03-3, School of Computer Science, University of Birmingham, October 2003.

[20] Viezzer, M.., Nieuwenhuis, C. M. "Learning affordance concepts: some seminal ideas", In *IJCAI05 Workshop on Modeling Natural Action Selection*.