The Virtual Chef
An Animated Pedagogical Agent

Sara Passos
IST - Technical University of Lisbon
Av. Prof. Cavaco Silva
Tagus Park
2780-990 Porto Salvo, Portugal
sara.passos@tagus.ist.utl.pt

ABSTRACT
Pedagogical agents are becoming more and more a part of our daily lives. They can assist and collaborate with students, playing a very important role in learning environments. These agents not only can provide proper advices but also perform practical demonstrations, or even replacing the human teacher. The main goal of this work is to build a model of an animated pedagogical agent, in particular, a virtual chef, and to find the more adequate tutoring strategies a cook should employ and why.
In order to achieve that, a study of several animated pedagogical agents was made, where the several tutoring strategies were analyzed
We believed that, if using the appropriate tutoring strategies, the virtual chef may improve the learning process, making it more enjoyable.

Keywords
Pedagogical agents, tutoring strategies, learning environment, learning process

1. INTRODUCTION
Pedagogical agents play a very important role in learning environments [6]. In particular, embodied pedagogical agents can also assume the role of "partners": they can perform multi-person tasks, play the role of a missing member. Embodied pedagogical agents are in constant growth and offer a great promise to interactive learning environments. It is enormous the evolution that this area has suffered over the last years, where those textual dialogue systems on a terminal have been successfully replaced with all those complex lifelike animated agents.
However, there is still a lot of work to be done. Despite these first efforts that resulted in some amazing pedagogical agents, there is still room for improvements in this area. There are several domains which have not been covered yet, and whenever that happens may cause a great impact.
"Cooking" is one of these unexplored domains.
Nowadays, there are several sites on the Internet with an enormous variety of recipes and tips from all over the world. There are also appearing several new cooking programs on TV. There are some chefs that became famous because of their programs, such as Jamie Oliver. These cooking programs are having a lot of success maybe because it is much more interesting and enjoyable to see a recipe being done, instead of reading it in a book. Besides, it is much easier to memorize something if we actually see it.
Having this context, this work tries to find out the main strategies used on the most conceived examples of animated pedagogical agents, and which of them can be considered to be the most effective in their respective domains. Taking these strategies into account, this work aims to create an animated pedagogical agent, a virtual chef, that due to its tutoring strategies, is capable of providing a better learning process.
In the next section, some the tutoring strategies employed in the area of pedagogical agents, in several domains.
Then, section 3 will described the basic architecture of the system and the main aspects that should be taken into account when conceiving a virtual chef, including the tutoring strategies that should be followed and why.
Afterwards, section 4 explains the evaluation the system that our system was subject and present some preliminary results.
Finally, section 5 presents some conclusions and future work.

2. RELATED WORK
Pedagogical agents for knowledge-based learning can provide customized support and advice for student's problem solving. They have two essential advantages: increasing the communication between students and computers and increasing the computer's ability to motivate students [2]. The important features of a pedagogical agent can be divided in two groups: tutoring strategies and elements that aid learning.

2.1 Tutoring strategies
Pedagogical agents appeared with the goal of facilitating and improving the learning process and experience. In some contexts, they will substitute human teachers and, therefore, must be as effective as possible, trying always to be distinguished as "good teachers".
In order to succeed as a tutor, the agents should implement some of these important features: Interactive Demonstrations, Navigational Guidance, Virtual Teammates and Adaptive Pedagogical Interaction.
When an agent inhabits a virtual world, it is possible to perform "live" demonstrations and, therefore, explain how to execute a certain physical task. The demonstrations can also include spoken explanations, so the agent can describe its actions while actually performing it.
When the environment is complex, students frequently get disoriented and, therefore, navigational guidance is very useful. Using this feature, the agent can guide the students, leading them around, showing the relevant objects and places and how to get there, preventing them of getting lost.
Virtual Teammates are used hen a certain task is very complex and requires different members to perform coordinated tasks. For example, teamwork is critical in an emergency room or in a battlefield
Pedagogical agents must be as dynamic and adaptive as possible instead of being pre-planned and sequential. It is useful if they have the ability to generate explanations, ask questions, track the learner's skill levels or generate dynamic and appropriate advices.

2.2 Elements that aid learning

The instruction strategies previously defined aren't enough for a proper and effective learning. There are some other elements that enhance learning. These elements have the main purpose of capturing and focalizing student's attention, providing appropriate feedback and encouraging and motivating the students. They are: Attentional Guides, Feedback, Conversation Signals and Conveying and Eliciting Emotion.

In a learning environment, it is very important to maintain the student always focused on the correct object or task. Agents may use many strategies to guide the user's attention. In animated agents the most common strategy is using gaze and gestures as attentional guides. Pointing at objects while explaining what they're for, looking directly at an object just before manipulating it or looking at students waiting for them to speak or to execute a task, are some examples of attentional guides.

One of the most important characteristic in a tutor is the ability to provide feedback to students' actions. With feedback, students may easily understand when they are executing correctly a task or making some mistakes. The agent can provide verbal or nonverbal feedback. Nonverbal feedback can prevent the user from committing a mistake and, sometimes, it is more appropriate than a verbal feedback, once is less perturbing.

Conversational signals are very useful in face-to-face conversations, helping it to flow more easily and smoothly. The agent may use pitch accent to indicate the degree and type of salience of words and phrases in an utterance. It is possible to highlight these utterance elements by accompanying pitch accents with short movements of eyebrows or head or blink of the eyes [1]. The agent can also combine nonverbal signals (such as eye contact) with spoken utterances to regulate turn talking in mixed-initiative dialogue.

Motivation is one of the key elements in learning. An emotive pedagogical agent can motivate students, conveying enthusiasm for the subject matter [5]. Animated agents can improve student's learning in various ways [2]: (1) if an agent appears to care about the student's progress, may encourage the student to care about his own progress; (2) when an agent is able to convey enthusiasm for the subject matter, then he may potentiate the same levels of enthusiasm in the learner and, finally (3) if an agent has a rich and interesting personality, the learning may become more fun.

All the features described above are considered the most important ones in order to provide a learning process as effective and as good as possible.

It's important to mention that none of the existing pedagogical agents implement all features described. Each characteristic, per se, is a real asset. Combining them differently can provide great and interesting results.

3. THE SYSTEM

Using these tutoring strategies as guidelines, and since cooking is still an unexplored domain in pedagogical agents, an interactive system was idealized, a virtual chef, which tries to contemplate the characteristics of real cooks, taking always into account the strategies studied before.

Figure 1, described the basic architecture of the virtual agent.

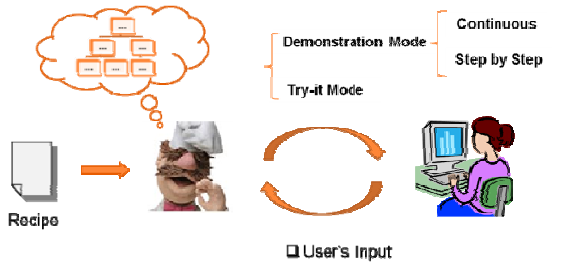

Figure 1 – Architecture of the virtual chef

This virtual chef is an animated pedagogical agent, focused on the cooking domain. First of all, he receives a recipe, and, based on it, can perform two distinct tasks: execute a demonstration or observe and guide a user performing that recipe.

In the first one, the agent performs a demonstration of the recipe using gestures and textual explanations, being always aware of the user's feedback. The user may ask agent to repeat some step whenever he needs and the agent must interrupt his demonstration and step backwards in his explanation.

In the second interaction mode the agent observes the user performing a recipe. Although the agent doesn't assume an active posture in this interaction mode, he must always be aware of user's input as well. According to the user's performance the agent may need to help and guide him in order to complete his task. This help may be reflected in some hints or tips or even another demonstration of that particular step, which is performed using, again, gestures and textual explanations.

So, in both interaction modes, the agent is constantly receiving feedback from the users which directly influences his reactions, choices, feedback and demonstrations.

The following sections will describe the main requirements in order to build a virtual chef. First, the structure of the recipe will be described as well as an overview of the concepts behind it. Afterwards, both interaction modes will be depicted where the strategies adopted in each of them are explained. Finally are made some remarks on the body module, where language and gestures generation are described.

3.1 The recipe

The recipe assumes a very important role in the agent's behaviour, since all the actions and gestures performed by the agent are based on it. The recipe's basic structure is illustrated on Figure 2, below.

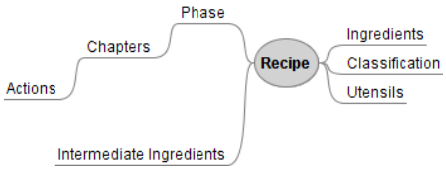

Figure 2 – Recipe's structure

Ingredients represent each ingredient of the recipe. These are the basic ingredients and, comparing with a normal recipe, they can be for example, sugar, wheat flour, eggs, among others.

Intermediate ingredients are the ones that result from the combination of other ingredients. For example when we mix sugar and yolks it will resolve into a cream.

Concerning classification, this element keeps some characteristics of the recipe such as duration or difficulty

Each recipe may have one or two phases: preparation and/or confection. The preparation phase occurs before even starting the recipe while the confection one is when the chef starts to cooking the ingredients.
Each phase contains a set of chapters which group a logic sequence of actions.
Concerning the action, this is the basic concept of the recipe. It contains information relative to each step of the recipe such as the description, the kind of the action (stir, melt …) or the ingredients involved.
Finally, the utensils are needed in every recipe and are a precious help to perform certain actions.

The structure of this recipe was conceived based on an ontology [9] made especially for the cooking domain.
The input file is received in a XML format and all information is read and saved for future use by the agent. The agent needs, then, to build a tree with the path he will need to follow all over the demonstration. To build this tree it was necessary to create an effective data structure where the runtime access was quick and simple, so the agent could have access the information promptly. In particular, the agent should be able to return to the correct path whenever interrupted by the user and needed to step backwards to repeat an action.

Considering Figure 1, the user's may interact the virtual chef by using two different modes: Demonstration and Try-it.
The following lines will describe the main strategies used behind each one.

3.2 Recipe Reasoning Engine

The entire recipe's information needs to be organized for a future use by the agent. He is responsible for choosing the right path to guide his demonstration as well as adapting it depending on the user's "performance".
All the information processing is done by a Recipe Reasoning Engine module. This module has the goal of building a tree of actions based on the recipe chosen by the user. To build this tree, it was created an algorithm, which is depicted below, that evaluates the several actions and orders them depending on their relations with each other:

- The root element of the tree is the last action of the recipe. In other words, the algorithm starts with the final result, for example, a Chocolate Mousse;
- Afterwards, we need to find the ingredients needed to achieve that result;
- Next, it is necessary to find the action that provides those ingredients (i.e. find the action where the final ingredients are the same as the ingredients needed by the last node).
- Add the action found as a new children node;
- Repeat these steps recursively for each node.
- The algorithm ends when the ingredients needed of a node are the basic ones

Figure 3 presents an example of a tree at the end of the algorithm.

3.3 Demonstration mode

For each mode were defined which could be the appropriate tutoring strategy as well as the elements that aid learning. The tutoring strategies chosen for this mode of Virtual Chef were: Interactive Demonstrations and Adaptive Pedagogical Interaction.

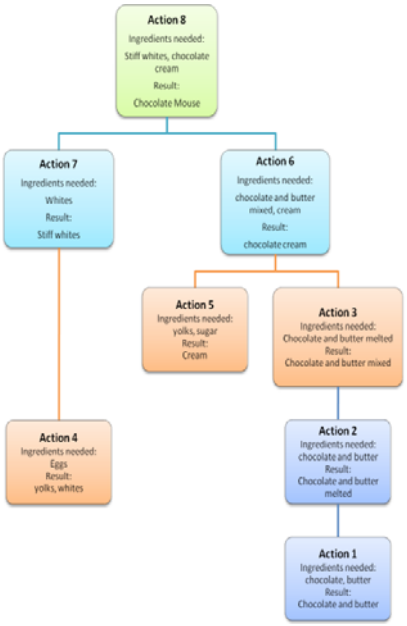

Figure 3 – Tree of actions at the end of the algorithm

Considering the purpose of this interaction mode it is quite direct the choice of using the interactive demonstrations strategy. Virtual Chef needs to perform practical demonstrations. A kitchen may be considered a complex environment with all its appliances, utensils and ingredients and, therefore, it is easier to follow and people tend to pay more attention to demonstrations because they can actually see all steps being done.
In order to execute the recipe, the cook must have a tree of actions that represents the path he needs to follow. However, this path shouldn't be static. The agent may need to change his route and should permit the user to interrupt his explanation whenever needed. Another important feature is that the cook should adapt his explanation according to the student's level. Depending on the user's knowledge the agent may need to provide more detailed explanations. Figure 4 illustrates the Virtual Chef performing a demonstration.

Figure 4 – Virtual Chef performing a demonstration

Making a recipe it's a very complex task. Not only due to the diversity of ingredients and possible actions but also due to the large number of steps a recipe may have. Having this, it is important that the Virtual Chef can collaborate and assist the users when needed, adapting his explanations and guiding his

demonstrations according to the user's level and providing some adequate tips or hints whenever necessary.
To support this strategy two "sub-modes" or levels were created The first sub-mode is the "normal" one where the agent explains the recipe in a continuous mode. During this demonstration, the user has the opportunity to interrupt the agent whenever he doesn't understand some step of the recipe and ask him to repeat. Then, the agent tries to explain again, however in a different way. There are two difficulty levels visible on the type of explanation. The default one is the medium level where the agent explains the essential. When the user shows some difficulty in understanding the recipe the system switches to the easy mode where the explanations are more detailed.

Concerning elements that aid learning it was chosen Attentional Guides.
By using these elements, the agent can capture more easily the students' attention where deictic gesture has a very important role. Therefore, it is important that the Virtual Chef includes attentional guides so the user can stay focused on the recipe's execution. The agent should be able to point at an ingredient or utensil unambiguously. Thus, the user may easily follow the agent's explanation with the interactive demonstration.
To support Attentional Guides, virtual chef implements an important feature. Parallel to the demonstration cycle there is a panel with several icons which represent the actions and ingredients of the recipe. These icons become highlighted whenever the respective ingredients/ or actions are referred in the explanation. These icons help the user to memorize which ingredients and actions were used as well as their respective combinations so he can identify them more easily.

3.4 Try-it Mode

This is the second possible interaction with the Virtual Chef: observing a user executing a recipe.
The tutoring strategy chosen for this second mode of Virtual Chef was Adaptive Pedagogical Interaction. Although this strategy has already appeared in Demonstration Mode, in this one it is used and applied in different conditions.
In this interaction mode, the cook presents a more passive behaviour. But "passive" doesn't mean indifferent. Although the agent doesn't have an active role, he must always be aware of all student's movements. This may be the difference between a good and a bad tutor. The fact of observing the user's actions may prevent several mistakes. The cook must be able to "feel" whenever the user is in trouble. When this happens, he should be capable of providing appropriate tips or hints, i.e. the tips must be directed to the problem the user is facing. If necessary, the cook must intervene and explain again the step or action to the student.
To support this strategy, the Virtual Chef creates his tips based on the type of mistake. For instance, if the user chooses the correct ingredients but the wrong action, the agent will give a tip concerning the action. Moreover, different types of help are triggered off also depending on the number of mistakes of the user as explained next:

- When the user makes one mistake the agent just informs that he made a wrong decision and incentives him to try again;
- At the second mistake, the agent gives him a hint;
- If the user still makes a wrong choice for the third consecutive time, the agent performs a "live demo" of that specific action just as in Demonstration Mode.

When concerning the elements that aid learning, the virtual chef implements Feedback and Conveying and Eliciting Emotion.
The Feedback strategy should be included in all systems where the user can actually interact with it, independently the domain chosen. Therefore, it was important that Virtual Chef implemented this strategy. When the user is executing a recipe he may commit mistakes and, by providing appropriate feedback, the cook may even prevent the user from committing some more mistakes in the future.
Virtual Chef includes feedback in several situations. When the user is performing an action he should always have the information if he is performing it in a correct way or not. It is important that the hints or tips are integrated in the feedback given, so the agent can not only inform the mistake but also help the user to recover from it. Figure 5 illustrates the virtual agent giving a feedback to a user.

Figure 5 – Agent providing feedback to the user

Finally, motivation is one of the key elements in learning environments. Concerning the Virtual Chef, it is important he includes this strategy. Due to the large number of steps a recipe may have, the user may become bored, so a constant and pleasant encouragement can make a difference in the user's learning process.
Feedback should always include some encouraging and motivating comments, independently if it is a good or a bad one. Thus, when the feedback is good and the agent still motivates the user, he gets even more encouraged and confident. But, when the feedback is concerning a mistake, an encouragement always feels nice and incentives the user not to give up.

3.5 The character

In order to perform all the strategies described on the previous sections, the agent has the constant support of two important features: utterances and gestures.
Utterances are generated by a Language Engine [11] and it only produces sentences that will be written on the screen. On its turn, gestures are produced by an Animation Engine. Both elements are used in parallel having, for each action, simultaneously an utterance and a gesture.

The animation engine is the responsible for all the agent's gestures.
To perform the gestures, it is used a module which is a part of a master thesis work called *Gesticulation in Virtual Humans* [10].
At the beginning, the system receives a XML file with all the gestures definitions, which will create the agent's gestures library. This file information such as the name of the gestures,

the position of the arm, the orientation of the hand, the hand and arm to move and the duration of the gesture.
This information is used in real time by the Virtual Chef. Whenever the agent wants to perform a certain action, he accesses his gesture's library, and executes the gestures with the information retrieved.

On the other hand, while the agent is performing a certain gesture, there are being generated utterances that represent his textual explanations. These utterances are generated by an external module, where we only have to define the structure of each sentence and provide the respective variables. Each sentence is built according to the type of speech (feedback or explanation) and uses information of the action.
Thus, it is possible the generate utterances as much dynamic as possible always contextualized with the current action.

4. EVALUATION

To evaluate the proposed solution some tests were performed. A total of 12 participants took part in this experiment (7males and 5 females). All of them were students on the IST – Technical University of Lisbon and their ages were comprehended between 20 and 27 years old. All the participants were native Portuguese and the only criterion followed in the selection of them was to choose people that have never interacted with Virtual Chef before.
The experiment consisted in two different phases: in the first one the agent performed a demonstration of a recipe. This demonstration could be performed in the continuous mode or in the step-by-step mode.
The second phase consisted in the user performing the recipe that had just learned. Thus, measuring some aspects such as number of error committed or how many times the user asked for help, it was possible to evaluate what was the demonstration mode that revealed to be more useful in the learning process.
By evaluating the user's learning process, we will, consequently, be analyzing if the agent's tutoring strategies are being effective and helpful. There are two ways of assessing this feature. The first and immediate one is measuring quantitative aspects such as time or mistakes the user committed while performing a recipe after seeing the virtual chef's demonstration. In this case, it was measured the following aspects: number of mistakes committed by the user, number of tips given by the agent, number of demonstrations given by the agent, number of times the user asked the agent to repeat, number of times the user exceeded the expected time, number of times the user asked for help ("Help" button)
The second way of evaluating this aspect was by using a questionnaire. At the end of the tests, it was distributed a questionnaire trying to assess the impact of the virtual agent and it's strategies on the user's interaction.

The results of this evaluation were obtained through our questionnaire results as well as measuring the metrics previously described and observing the user's reaction constantly.

With the help of the questionnaire we could conclude that the assistance provided by the agent turned out to be adequate to the difficulties the users were facing while performing the recipe. Concerning the user's enjoyment users thought the agent was very pleasant and in, some cases, even funny. The reaction was quite good. However the introduction of voice and/or sounds would provide an even better experience to most of the users.

Concerning the quantitative results, we observed that the participants that were submitted to the continuous demonstration mode presented better results than the ones that interacted with the step-by-step one. Although some considerable mistakes appeared on both modes, this cannot be a negative aspect of the Virtual Chef since the system was new to all the users and the recipe had some considerable steps. Mistakes were expected, but the point to prove was that the users could recover their mistakes with the help provided by the agent and therefore, complete the task they began.
Another aspect that was possible to conclude was that the step by step mode turned out to be a "two way road" because, in one way the fact of the speed of the explanations being according to the user's will, can help them to better understand each step but, in the other hand, the possibility of distraction is much higher than the in other mode. Despite this, the inclusion of this mode in the system still seems appropriate because these distractions didn't affect all the users and some of them even provided good results.
However, some tests can still be made. These were preliminary tests and, in order to conclude if the presence of the Virtual Chef really makes a positive influence on the user's learning process, it could be performed a similar test to the ones made, but wouldn't have a demonstration phase. The user only had to read the recipe in a paper and, then, executes it interacting with the Try-it! Mode. By assessing the number of mistakes committed and all the metrics used, it was possible to conclude if the presence (demonstration) of the agent had really made a difference.

5. CONCLUSION

The work presented in this document pretends to provide a contribution in animated pedagogical agents. They are becoming more a more an important tool in learning environments. They can provide proper advices and, in some cases, perform practical demonstrations; they can even replace a human teacher. These embodied pedagogical agents, when lifelike and believable, can turn the learning process more effective, motivating and enjoyable.
This work pretended to idealize and build an animated pedagogical agent, in this case, a virtual chef. For that, it was made a deep study on tutoring strategies so it was possible to infer what could be that best tutoring strategies to build a virtual chef.
After implementing a system that could support the model idealized, a preliminary evaluation was performed. The first measure to evaluate was the user's learning process (by analyzing the effectiveness of the tutoring strategies chosen) and second one was the user's enjoyment. Although the number of tests were limited, it was possible to take some positive conclusions. The majority of the participants liked to interact with the agent and even declared that they had learned how to perform the recipe with the agent's help. However, some other experiments still needs to be done so it can be possible to assume that this virtual chef can really be an asset to the user's learning process.

6. FUTURE WORK

As described before, it is essential that a cook performs interactive demonstrations. A kitchen may be considered a complex environment with all the utensils and ingredients present in it. It was interesting to add some movement to the Virtual Chef besides the gestures. The cook should be able to walk through the environment during his explanations, such as moving towards the oven whenever is necessary to melt or heat something. This feature can increase the agent's

believability and therefore can enhance the user's learning and enjoyment.

Feedback is one of the agent's strongest aspects. However there is one important feature that could be introduced: facial expression. Virtual Chef only uses nonverbal feedback through gestures that represent his agreement or disagreement. By entering facials expressions, the feedback would become much more complete and less intrusive. For instance, the agent could prevent some mistakes but just making a sign with his eyes showing the user he might be following the wrong way.

Another aspect that should be taken in consideration is the synthesized speech. Some users explained it was hard to follow the gestures and read the sentences at the same time. Therefore, people tend to focus only in one of them, generally the utterances. If the cook could speak, users would limit to listen to him and could pay all attention to the gestures which are an important guide. Besides, the fact of the application having sound would increase the enthusiasm of the user, instead of just looking at a silent demonstration.

Finally, it was interesting if the agent could actually manipulate the ingredients. This would give a more realistic impact of the system as well as turning it much more interesting and pleasant to interact with.

This dissertation tried to provide a contribution to the area of pedagogical agents, in particular, pedagogical agents in the cooking domain. This work meant to be innovative, since it doesn't exist any other work using this domain. Moreover, by trying to discover what where the strategies that best suit a virtual chef, a good study on tutoring strategies was made. This can be useful not only to a virtual chef. By knowing and understanding each of the strategies it is easier to apply them correctly in any domain.

Although this Virtual Chef may not be as complete as desired, it can be a great point to start entering some new areas in pedagogical agents.

7. REFERENCES

[1] Cassell, J., et al. "Animated conversation: rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents." Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques. ACM, 1994. 413-420.

[2] Johnson, W. L., J. W. Rickel, e J. C. Lester. "Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments." International Journal of Artificial Intelligence in Education, 2000: 47-78.

[3] Lester, J. C., e B. A. Stone. "Increasing Believability in Animated Pedagogical Agents." 1997. 16-21.

[4] Lester, J. C., J. L. Voerman, S. G. Towns, e C. B. Callaway. "Deictic Believability: Coordinated Gesture, Locomotion, and Speech in Lifelike Pedagogical Agents." Applied Artificial Intelligence, 1999: 383-414.

[5] Person, N. K., e A. C. Graesser. "Pedagogical Agents and Tutors." In Encyclopedia of Education, de J. W. Guthrie. New York: Macmillan.

[6] Rickel, J., e W. L. Johnson. "Integrating Pedagogical Capabilities in a Virtual Environment Agents." Proc. of First International Conference on Autonomus Agents. ACM, 1997. 30-38.

[7] Rickel, J., e W. L. Johnson. "Task-Oriented Collaboration with Embodied Agents in Virtual Worlds." In Embodied Conversational Agents, de J. Cassell, J. Sullivan e S. Prevost. Boston: MIT Press, 2000.

[8] Rist, T.; André, E.; Müller, J.; "Adding Animated Presentation Agents to the Interface." In Intelligent User Interfaces, 1997. 79-86.

[9] Ribeiro, R., Baptista, F., Pardal, J., Mamede, N., & Pinto, H. (2006). "Cooking an Ontology". AIMSA, (pp. 213-221).

[10] Melo, C. "Gesticulation Expression in Virtual Humans", Master Thesis, Instituto Superior Técnico, Lisbon, 2006.

[11] Dias, J., Ho, S., Azenha, B., Prada, R., Enz, S., Zoll, C., Kriegel, M., Paiva, A., Dautenhahn, K., Ayllet R., Gomes, M. R., Ferreira, C.. "D4.1.1-D6.2.1 Actor with Attitude and Component Specification for Computational Autobiographic Memory", (2007), ST-4-027656-STP eCircus