Where Did I Put My Glasses? Determining Trustfulness of Records in Episodic Memory by Means of an Associative Network

Cyril Brom[1], Klára Pešková[1], and Jiří Lukavský[2]

[1] Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic
[2] Institute of Psychology, Academy of Sciences, Prague, Czech Republic
brom@ksvi.mff.cuni.cz

Abstract. Episodic memory represents personal history of an entity. Human-like agents with *a full episodic memory* are able to reconstruct their personal stories to a large extent. Since these agents typically live in dynamic environments that change beyond their capabilities, their memory must cope with determining trustfulness of memory records. In this paper, we propose an associative network addressing this issue with regard to records about objects an agent met during its live. The network is presently being implemented into our case-study human-like agent with a full episodic memory.

1 Introduction

From the psychological point of view, *episodic memory* [17] represents personal history of an entity. Episodic memories are related to particular places and moments, and are connected to subjective feelings and current goals. *Human-like agents* (or *virtual humans*) are typically thought of as software components *imitating* behaviour of a human in a virtual world that are equipped with a *virtual body* graphically visualised. The important feature of a virtual human is that it is designed to imitate behaviour in a *believable* manner, but not necessarily psychologically plausibly. These agents inhabit artificial worlds, be it 2D or 3D, in commercial computer games, serious games, virtual storytelling applications, military simulations and various other applications (for a review, see [15]). These worlds are typically dynamic and unpredictable, and the user interaction is often allowed.

Human-like agents with an *ad hoc* episodic memory [5] are able to store only the events specified inside an agent's script or a reactive plan in a hardwired fashion. This kind of episodic memory is almost always present in current human-like agents for it is essential for action selection purposes. In contrast, a *full episodic memory* stores more or less everything happing around the agent in a general manner tagged with the agent's own relevance estimation. This presents actually a form of a life-long, autobiographic memory [11], which is often absent in current human-like agents. However, there is a growing need of agents with this kind of memory in the fields of narrative storytelling systems and role-playing computer games, as discussed e.g. in [7]. For example, these agents are able to reconstruct their personal history, which increases their believability. Imagine that while you are playing a role-playing game, you come to a medieval town, enter a magic shop and ask the computer driven seller:

F. Almeida e Costa et al. (Eds.): ECAL 2007, LNAI 4648, pp. 243 – 252, 2007.

Hey, I am the representative of the king Hromburac Pekelny, and you please tell me, what were you doing during the last week? And, please, summarise it in three sentences.

Well – answers the seller – I was in my shop every day except of Friday, selling my magical stuff. In the evenings, I enjoyed with my friends. Nothing interesting happened, except of Wednesday when a filthy usurer came to rob me.

Yet, a full episodic memory agent will allow you to ask further, for example:

Ok, that sounds interesting. Please, summarise the same now in 15 sentences. or
Focus please on the filthy usurer. or
Tell me more about Saturday evening.[1]

We have developed a prototype of such an agent with a full episodic memory [5] (Fig. 1) as a part of our on-going work on an educational storytelling game [3]. The memory stores the course of activities the agent performed, the objects the agent used, and also the reasons why the actions were performed (to some extent). It also forgets unimportant episodes as a time passes. It is widely optimised for storing and retrieval.

During development, we have stumbled on several inherent episodic memory issues, some of which remain to be solved. One of these issues is the problem with estimating trustfulness of the memory records that relate to objects. As the world is dynamic, an object may move beyond the agents capabilities (e.g. by another agent). Hence, after a couple of days in the virtual world, the agent may have tens of records about the same object. Which is the most trustful one?

In this paper, we address the issue of trustfullness by proposing an associative network aimed at coping with this problem. Implementation of this network is our current work-in-progress. Although the network is aimed as a component for our episodic memory employed in human-like agents, it can be also used *as is* in other artificial life agents.

Section 2 overviews our previous work concerning episodic memory and Section 3 proposes the associative network. In Section 4, an overview of related work is given.

2 Overview of Previous Work

This section details requirement on the full episodic memory architecture we had, and briefly reviews the overall architecture of the agent and its episodic memory.

As our goal was to *imitate* human-like episodic memory, we needed agents to remember only what real humans would remember, and forget in similar way and extent as real humans would do. Unfortunately, there is no thorough functional description of episodic memory of humans from which we could derive our model. Thus, we were forced to derive the requirements only from case-studies of forensic psychology like John Dean's testimony study [10] and from our own phenomenological experience:

[1] Of course, apart from the episodic memory, the agent must be equipped with a *linguistic module* allowing for transferring the outcome of the memory to syntactically correct sentences, and the player's question into a memory query.

1. The memory should cope with complex tasks that require manipulation with several objects and apparently require human-level cognitive abilities, like cooking or shopping, for they can be performed by a human-like agent. Such tasks tpically have hierarchical nature—they can be logically decomposed to sub-tasks, which can be divided to yet smaller tasks, until some atomic actions are reached.
2. The memory has to store and reconstruct personal situations: a) what an agent performed with which objects and why, and b) who was seen by the agent and what this second agent performed (we remark, that presently only (a) is implemented, (b) is another work-in-progress). Time information, at least approximate, shall be also stored. The memory should be able to provide information like “where is something?”, “when did the agent see *x*?”, “what did the agent do from *y* to *z*?”, “why did you do *a*?”[2], and reconstruct stories like the above-mentioned. The memory is expected to reconstruct information *on demand* (i.e. when a user asks) rather than automatically based on an environmental clue.
3. The memory should not store all available information; neither external, nor internal. In particular neither all objects located around the agent, nor all activities of other agents should be stored. The objects the agent uses shall be stored more often than the objects not used, but only seen. Only important actions performed by other agents should be stored. Not all internal state should be stored, but only information about motivations relevant to current goals (i.e. I would not remember I am slightly thirsty, if I am dying hungry at the same time and my goal is food and only food). Generally, the activities/objects to be remembered should be chosen based on their relevance to the agent, on their general attractiveness, and on the agent’s attentional and emotional state.
4. The memory should operate in a large time scale. As time passes, the unimportant details should be forgotten, until only a “gist” of what happened is kept. Different episodes should be forgotten in different speeds based on their importance and emotional relevance. Several similar episodes can be eventually merged together.
5. Coherence shall be maintained, in particular if there are two contradictory records in the memory, i.e. an object *x* has been seen both at *a* and *b*, one of them must be marked as more trustful. This last point is actually the scope of this paper.

Previous memory model. As mentioned, the associative network is primarily intended as a plug-in for the full episodic memory agent we developed previously. For brevity, we only sketch this previous model here. Full details can be found in [5].

Architecture of our agent is depicted in Fig. 2. Perhaps its only distinction from a classical symbolic architecture is the episodic memory. All the parts except of the linguistic module are implemented, emotional module being implemented separately [2], and its integration is in progress.

The agent is driven by hierarchical reactive planning with behaviour represented by AND-OR trees [4]. Basically, the AND-OR tree metaphor works with abstract *goals* representing what shall be achieved, and *tasks* representing how to achieve the goals. Typically, every goal can be accomplished by several tasks, while every task

2 Causal inference, i.e. “why”, is complicated, but an agent can answer these questions to a limited extend for it has some basic understanding of its actions owing to its behavioural representation. It has a goal and it picks the most suitable action for it in every moment.

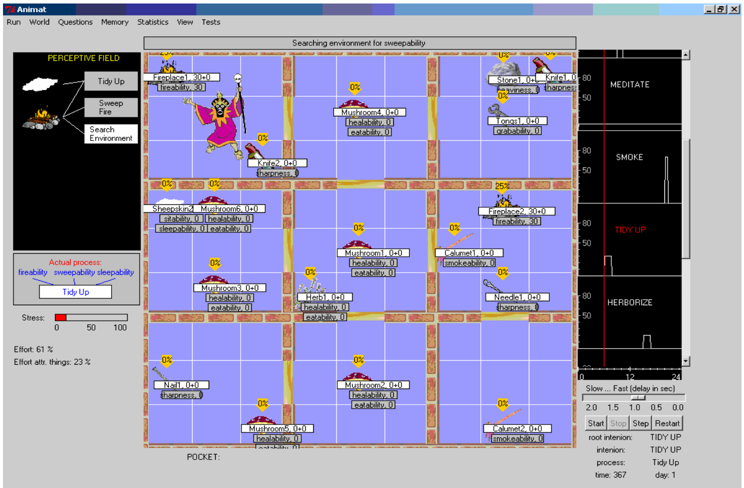

Fig. 1. A screen-shot from the prototype. Left: the short term memory. Middle: the 9x9 grid world of the shaman agent (the left upper part). Right: activation of the top-level goals.

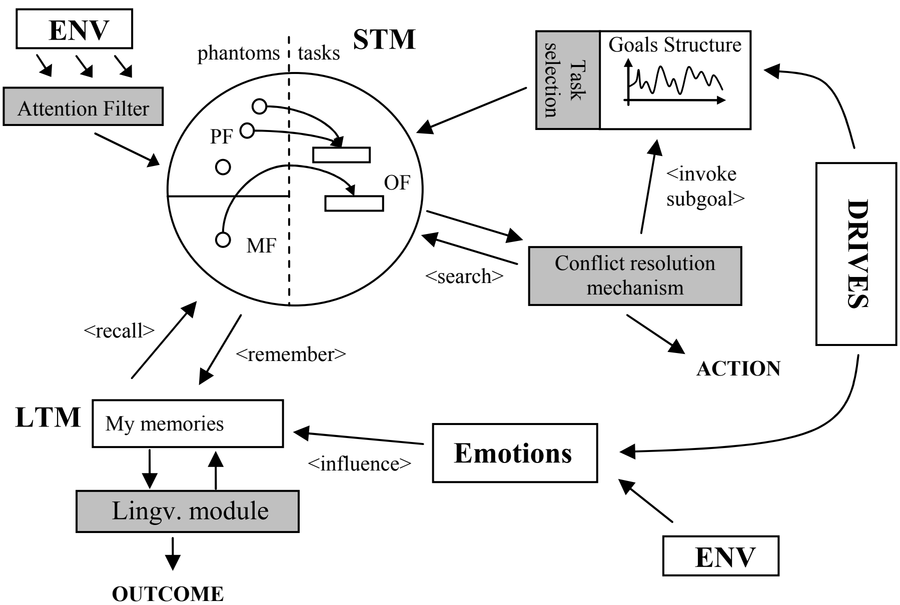

Fig. 2. The overall architecture of our agent. PF – phantoms of the STM. OF – own tasks of the agent. MF – records retrieved from the LTM. Emotional module is presently implemented separately (the emotional influence being specified manually), linguistic module remains unimplemented.

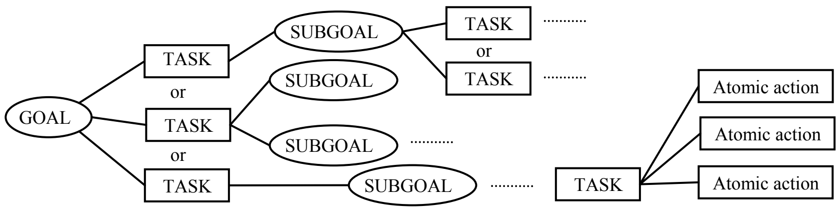

Fig. 3. Agent's behaviour is represented using AND-OR trees.

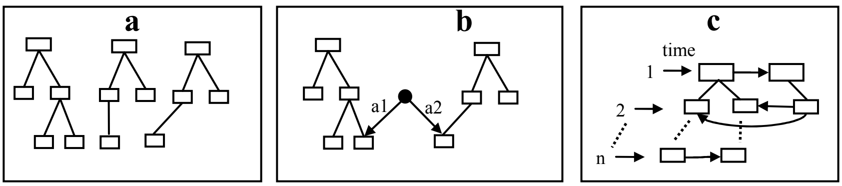

Fig. 4. a. A fixed arrangement of the long-term episodic memory, each box represents a task. 4b. Storage of a phantom of an object. This object can be used as a resource for two tasks (pointers *a1*, *a2*). 4c. The tasks are being sorted by time pointers during storing.

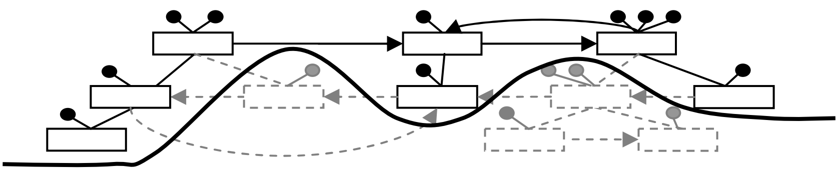

Fig. 5. The LTM forgetting schematically depicted.

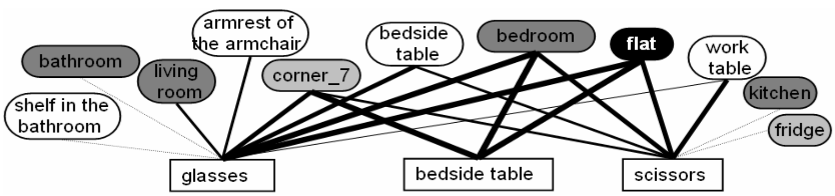

Fig. 6. A part of the associative network depicted. Intensity of grey denotes the level of space abstraction. The strength of the edges denotes the weight of the associations.

can be achieved by adopting some sub-goals. The agent needs to perform only one task to achieve a goal, provided there is no failure (hence, OR nodes), but to fulfill all sub-goals to solve a task (hence, AND nodes; Fig. 3). The tasks that cannot be further decomposed are *atomic actions*, i.e. primitives changing world-state. Every task may need several *resources* to be performed, i.e. objects (e.g. hammering is possible only with a nail and a hammer). Every top-level goal has its *activity level* based on drives, external events, and a schedule—the activity level is the mechanism of competing among the top-level goals. This competition takes place within the *goal structure*, which also stores the AND-OR trees. The winning goal chooses the most appropriate task (e.g. "to eat" goal can chose "take something from the fridge") and passes its template to the *short-term memory* (STM). Apart from these templates, the STM holds templates of objects seen that passed through the *attentional filter*. The object templates are called *phantoms*. In fact, they are classical index-functional entities [1]. Owing to a decay mechanism, there can be up to about 8-10 such phantoms and task templates in the STM, which is roughly consistent with human data [9].

If a task needs an object that is not seen, i.e. there is no relevant phantom in the STM, *the long-term memory* (LTM) is queried, and the environmental search is initiated. This may lead to invoking of a new sub-goal, and eventually a sub-task.

The LTM is a tree-like structure comprising all the tasks the agent can perform (Fig. 4). During remembering, two types of entities are added into this structure: phantoms, and so-called *time pointers*. Each phantom in the LTM represents an object used as a resource of a task in a particular moment, or an object that was not used but attracted the agent's attention. The time pointers represent the course of events.

Perhaps the most important feature of the LTM is forgetting: the less important episodes are being "bitten out" from the bottom of the LTM (Fig. 5), i.e. the time pointers and phantoms of the "bitten out" parts of the LTM are being removed, typically during night as a "consolidation". The episode importance is presently determined by the age of the event (automatically), and by its emotional salience (manually).

Implementation and evaluation. We have prototyped the memory model and a case-study scenario in Python (Fig. 1, [5]). The scenario features a "shaman" agent living in a simplified grid-world for several days (3x3 rooms, each room 3x3 tiles). The shaman has about 15 different top-level goals, from which she focuses on about 10 goals a day. She is able to answer the following questions: Where is an object? Where was an object from *a* to *b*? What did you do from *a* to *b*? When did you do *x* lastly?

Several tests measuring the efficiency of the memory have been carried out [5, 13], revealing that the memory performs well concerning both the size (in terms of memory units), and effectiveness of the agent (in terms of time spent by searching for an object). Concerning the size, however, forgetting was essential in large time scales. Concerning the effectiveness, the problem with assessing trustfulness of the memory records rose in worlds with a high dynamic – the shaman often searched for an object that had been already relocated.

Problem revisited. When the LTM in our current model is searched for an object, more than one phantom of this object can be found: e.g. glasses on the table, next to the TV, at the armrest of the armchair etc. Which of these records shall be considered as the most trustful? Apparently, this depends on the history of the world; on the habits of the agent as well as on the habits of the other agents in case of an object being shared. This issue is remarkable because searching at wrong places undermines both the agent's believability and effectiveness. While the believability is related only to the domain of human-like agents, the effectiveness issue is more general. In the next section, we propose an associative network aiming to cope with this issue.

3 The Proposed Associative Network

This section proposes first generation of the associative network addressing the issue of determining trustfulness of different phantoms of the same object in the LTM. Its implementation and integration into the rest of the model (described in Sec. 2) is our current work. The section also discusses possibilities and limitations of the network, including proposing experiments to be carried out and directions of future research.

The network architecture is designed to cope with dynamic and complexity of human-like worlds. We, present-day humans, live in a world we typically describe in a hierarchical fashion; we think in rooms, houses, streets, places etc.[3] We tidy up objects: they are not placed randomly in our surroundings, but logically clustered according to our needs and cultural norms (e.g. a spade is typically not to be found in a kitchen). Objects we use have their specific dynamic: some are being moved every day but from a typical place, e.g. glasses from a bedside table; others are almost never being relocated, e.g. a fridge; yet others change their position so often that it does not make sense to remember their exact position. Consequently, during searching, we sometimes scrutinise a whole room, sometimes we inspect several very specific places, sometimes a stimulus-response mechanism is employed. For example, we may have the following rule for searching for glasses: to inspect the bedside table first, then to look at a work table, and eventually to scrutinise the whole flat, starting in the living room. In another situation, we may go to the bathroom to fill the washing-machine and then realise: aha! – it was moved to the kitchen a week ago.

Basically, this complexity and dynamic was intended to be seized in the network. We believe we have found a relatively simple architecture for it, yet fulfilling the requirements. The network comprises two kinds of nodes: *place-nodes* and *phantom-nodes*. A phantom-node is an index-functional entity representing a particular object. A place-node represents a place where an object can be found. By place we mean any logically coherent space abstraction, e.g. a bedside table, the place between this table and the bed, a bedroom, a corner of the bedroom, a flat, etc. A place-node can also represent an object, as in the case of the bedside table. The places *are* from different levels of abstraction, which are numbered; e.g. the bedside table is the level 1, the corner is the level 2, the living rooms is the level 3, the flat level 4 etc.

A phantom-node can be connected to several place-nodes. Such an *edge* represents a possible occurrence of this particular object at this particular place (Fig. 6). Every edge has its *weight* meaning the number of times the object was seen there. We remark that while in our previous LTM, a phantom for each occurrence of an object was stored in the LTM, in this network, only one phantom is stored for one object.

Formally, the network is a triple $<P, H, E>$, where P is the set of all place-nodes, H the set of all phantom-nodes, and $E = \{< x, w >\}$ is the set of weighted edges (where $x \in P \times H$, and $w \in \mathbf{N}$).

Learning mechanism. In addition to storing a phantom into the LTM from the previous model (Sec. 2, [5]), weights of the edges between the object's phantom and *all* the places where the object was found are increased by 1 (new places being added).

[3] The question of whether our predecessors, which did not live in rooms and houses, had hierarchical representations of their surroundings or not, is not discussed here for brevity.

E.g., if glasses are found on the bedside table, the edges to the bedside table, to the corner next to the bed, to the bedroom, and to the flat are strengthen. Consequently, the associations between a phantom and the places the object is typically being found are strengthen more often than the other edges, and the edges to more abstract places are also strengthen more often as the total number of places at a particular level of abstraction decreases with the number of the level (e.g. the glasses are almost always in the flat, but sometimes on the bedside table, and sometimes on the work table).

Retrieval. For every place the phantom has an edge with, a *size-normalised trustfulness* (SNT) is computed as w/a^b, where w is the weight, a is the number of the layer of abstraction of this place, and b a scaling factor. The places are sorted according to this value and the ones that are over a *treshold* are considered as places where the object may be located. The object is then searched for at these places in the order of the SNTs. The concrete places are looked at directly (e.g. a bedtime table), abstract place are to be inspected. The SNT is supposed to lead to a searching process that prefers concrete places over abstract ones, provided there are only a few concrete places. Otherwise, it is preferred to inspect the abstract place directly, because searching in the concrete places is likely to fail. The factor b is intended to be tuned during experimenting.

Issues. Some questions are to be investigated after we finish the implementation of the first-generation network. Most notably, these are (1) exceptions, (2) classes of objects, (3) edges weakening, (4) states of objects, and (5) moving objects.

1) Sometimes, it may be fruitful to remember exceptions; e.g. "the glasses are typically at the bedside table, but now, I have put them at the TV". Phantoms from the STM are actually preferred to LTM records, but the question remains whether the exceptions should be handled also in the LTM, e.g. when a task is interrupted, in which case all the phantoms of the objects used during the task decay from the STM.

2) In the described network, each object is stored separately, including objects that are being replaced frequently like food or newspapers. We suggest that to store only one phantom for the whole class of such objects is a better solution. We will focus on this implementation issue after evaluating the first generation of the network.

3) Should the connections in the network be also weakened? The trouble is that sometimes humans remember things over long periods, e.g. that the 1979 Bordeux is in the cellar, even if we did not see it for last 10 years. This suggests that edges can not be weakened in time straightforwardly. Instead, we propose (i) weakening an edge only if it is small comparing to the sum of the weights of the other edges originating from *this* phantom. Additionally, we propose (ii) intentional weakening ("I try to forget that the washing-machine was in the bathroom when it has been moved to the kitchen"), and (iii) weakening in case of a place that was suggested by the network, but where the object was not found (this is now achieved indirectly because the edges to the places where the object was found are strengthened). We plan to experiment with these three mechanisms.

4) One object can have different states. Sometimes, the states may be associated with places, e.g. full wine bottles are in the cellar, while empty bottles are next to the bin. Could this be mirrored by the network?

5) The memory is not designed for objects that move themselves (a dog). An entirely different memory would have to be used for this class of objects.

Another interesting issue, which is however out of scope of this paper and our current research, is investigation of relations between this network and mechanisms of natural episodic memory. One interesting phenomenon to look at is place cells [8].

4 Related Work

In the field of human-like agents, the issue of generic episodic memory has been almost untouched, since the agents typically need not store more than a few episodes for action selection purposes. A notable exception is the tutor agent Steve [16], who employs an episodic memory that allows him to explain himself after a given lesson, which lasts, however, only a couple of minutes. Another exception is the memory model for agents of ALOHA system [14], which exploited to a great advantage division of the memory to short-term one and long-term one, but unfortunately stores records only about objects and groups of objects (but not the tasks).

In robotics, Dodd [6] developed recently a general memory system for a humanoid robot ISAC, which included also episodic memory working with emotions. An interesting a-life example is the work of Ho et al. [7], who developed agents with various different episodic memories aiming to investigate how different types of memories improve survival of the agents. Though these memories fit well for their domain, they are relatively low-level from the point of view of human-like agents. They are not designed to cope with complex, hierarchical tasks for example.

In agent research, perhaps the most elaborate model of episodic memory has been developing by Nuxoll & Laird [12]. This model is intended as a robust, general-purpose architectural extension of Soar with much broader scope than the model of ours. This, however, means that our model can benefit from some domain-specific tricks, which may finally increase efficiency of the memory in our domain. In all cases, it will be interesting to compare this model with ours in future.

Typically, the environments presented with the abovementioned memories are not very dynamic. To our knowledge, none of the models addresses the issue of trustfulness of memory records to the extent as the proposed network does.

5 Conclusion

We have proposed an associative network that addresses the issue of assessing trustfulness of records about objects in the long-term episodic memory of a human-like agent. The network is intended to elevate believability of the agent, and its efficiency (measured in time spent by searching for an object) in highly dynamic human-like environments. The network is presently being implemented into our agent with a full episodic memory developed earlier [5].

The main idea behind the network is that a place where an object was found is represented by several nodes, each representing this place at a *different* level of abstraction. The network is relatively simple, which is, however, its advantage considering future rigorous analysing and extending.

We have introduced a mechanism for searching for an object by an agent equipped with the network. We have also discussed limitations of the network proposing some issues that are ahead. These includes, most notably, handling exceptions, weakening of the associations, and coping with classes of objects being frequently replaced. These issues will be addressed as a part on the work of the second generation of the network.

Acknowledgments. The authors would like to thank to Andrew Nuxoll for answering questions on his work. This research was partially supported by the Program "Information Society" under project 1ET100300517, and GA UK 351/2006/A-INF/ MFF.

References

1. Agre, P.E., Chapman, D.: Pengi: An Implementation of a Theory of Activity. In: Proceedings of the Sixth National Conference on Artificial Intelligence, Seatle, USA, pp. 196–201 (1987)
2. Bída M., Kadlec R., Brom. C.: Relevance of emotions for artificial beings. (Význam emocí pro umělé bytosti) (in Czech). In: Mind, intelligence and life (Myseľ, inteligencia a život), pp. 158-172. Vydavateľstvo STU, Bratislava (2007)
3. Brom, C., Abonyi, A.: Petri-Nets for Game Plot. In: Proceedings of AISB Artificial Intelligence and Simulation Behaviour Convention, Bristol, vol. 3, pp. 6–13 (2006)
4. Brom, C., Lukavský, J., Šerý, O., Poch, T., Šafrata, P.: Affordances and level-of-detail AI for virtual humans. In: Proceedings of Game Set and Match 2, The Netherlands, Delft (2006)
5. Brom, C., Pešková, K., Lukavský, J.: What does your actor remember? Towards characters with a full episodic memory. Tech. Report No. 2007/4 of the Dept. of Software Engineering, Charles University in Prague (2007)
6. Dodd, W.: The design of procedural, semantic, and episodic memory systems for a cognitive robot. Master thesis. Vanderbilt University, Nashville, Tennessee (2005)
7. Ho, W., Dautenhahn, K., Nehaniv, C.: Autobiographic Agents in Dynamic Virtual Environments - Performance Comparision for Different Memory Control Architectures. In: Proc. IEEE CEC, pp. 573–580. IEEE Computer Society Press, Los Alamitos (2005)
8. O'Keefe, J., Dostrovsky, J.: The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Research 34(1), 171–175 (1971)
9. Miller, T.: The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological review 63(2), 81–97 (1956)
10. Neisser, U.: John Dean's Memory: A case-study. Cognition 9, 1–22 (1981)
11. Nelson, K.: The psychological and social origins of autobiographical memory. Psychological Science 4(1), 7–14 (1993)
12. Nuxoll, A., Laird, J.E.: Extending Cognitive Architecture with Episodic Memory. In: Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, AAAI, Menlo Park (2007)
13. Pešková, K., Brom, C., Lukavský, J.: Episodic memory for virtual actors based on affordances (in Czech). In: Proc. of Cognition and Artificial Life, Slovak Republic (2007)
14. Peters, C., O'Sullivan, C.: A Memory Model for Autonomous Virtual Humans. In: Proceedings of Eurographics Ireland, pp. 21–26 (2002)
15. Prendinger, H., Ishizuka, M.: Life-like Characters: Tools, Affective Functions and Applications. Springer, Heidelberg (2003)
16. Rickel, J., Johnson, W.L.: Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control. App. Artificial Intelligence 13 (1999)
17. Tulving, E.: Elements of Episodic Memory. Claredon Press, Oxford (1983)