Proceedings of the Remembering Who We Are – Human Memory for Artificial Agents Symposium, Mei Yii Lim and Wan Ching Ho (Eds.), at the AISB 2010 convention, 29 March – 1 April 2010, De Montfort University, Leicester, UK

Timing in Episodic Memory: Virtual Characters in Action

Ondřej Burkert[1], Cyril Brom[1], Rudolf Kadlec[1], Jiří Lukavský[2]

Abstract. In many applications, for instance in role playing games, it is an advantage when "minds" of virtual characters feature an episodic memory system. This system can boost cognitive and learning capabilities of the characters as well as their ability to respond to player's questions. Recently, several special-purpose memory mechanisms for virtual characters have been published. We have been developing a more generic model, which incorporates hierarchically organized memory for events with gradual forgetting, a component reconstructing plausibly time when an event happened and a spatial memory for "what-where" information. One open question that has not been addressed yet in the context of virtual characters is how precisely should an episodic memory store timing information. To answer this question, we have conducted a study to investigate what time categories people use when asking time-cued questions. We hypothesized that humans prefer using fuzzy categories such as "morning" or "after lunch" rather than exact information and this hypothesis was confirmed. Here, we present the results of the study and overview the part of our memory model that is responsible for timing.

1 INTRODUCTION

Episodic memory [1, 2] is an umbrella term for memory systems operating with representations of personal history of an entity, a term stemming from neuro-psychology. Episodic memory traces are related to particular places and moments, and connected to subjective feelings and current goals. Fundamentally, the episodic memory is being distinguished from the semantic memory and the procedural memory. The former is conceived, more or less, as systems operating with general facts about the world as viewed from the objective perspective. The latter covers processes related to skill learning.

A believable virtual character, or intelligent virtual agent (IVA throughout) is a software agent [3] who imitates human or animal behaviour in a 2D or 3D virtual environment, who seems lifelike and whose actions make sense to the audience allowing them to suspend their disbelief providing convincing portrayal of the personality they expect or come to expect [4]. Recently, it has been argued that episodic memory is one of the key components contributing to the believability, at least for IVAs interacting with humans for more than a couple of minutes, because it allows the user to understand better the IVA's history, personality, and internal state: both actual state and past state [5, 6, 7, 8]. Examples of skills an IVA can posses that demand some facet of episodic memory includes, but are not limited to: a) general giving of information to the user based on the past history of the IVA's interaction with the virtual environment, b) remembering a course of interaction with a human, notably remembering the course of a dialog, c) debriefing after a lesson, d) learning, e) repetition detection and improved reasoning [5].

Many applications require IVAs to have some of these skills, including role-playing computer games, educational simulations, virtual companions, and applications featuring virtual tutors, implying the necessity of augmenting IVAs' minds with episodic memory systems [9]. Accordingly, several special-purpose episodic memory mechanisms have been presented recently [10, 11, 12, 13]. In our approach, we depart from these solutions in that we aim at a generic, reusable model integrating multiple interacting subsystems.

One of such subsystems is a module estimating time when an event happened. Without the ability of time events, human life would be dull. For instance, imagine yourself recalling morning that you brushed your teeth but not whether this happened a minute ago or three days ago. Virtual characters should be equipped with a similar ability. The proper timing is vital for more believable output of an IVA when attempting, for instance, to summarize last few days into several sentences highlighting the most important events to give a better picture to the user.

Human timing is not perfectly cleared up by psychology yet. Nevertheless, there is an abundance of theories trying to explain many phenomena linked with human timing as well as evidence from experiments. These materials are an important source of inspiration for a developer of such a memory component for timing.

After discussing several theories, Friedman [14] comes to the conclusion that humans use multiple techniques for dating past events. For instance, if you try to recall what you did on December 20, you would automatically convert the date to four days before Christmas and subsequently use this Christmas as an important landmark in the flow of time. You would remember that you were working the two days just before the Christmas resulting in the refinement of the time information to Saturday before the Christmas. Thus you would conclude that you probably went to buy gifts. This way of recall uses reconstructive nature of the memory when general semantic knowledge, knowledge of our lifestyle and personal episodes are combined to refine the recall. But it can also happen that you would answer that you do not know as it was too long ago and the date is too exact. Consider now another example: recalling when you visited your dentist last time. It is likely that you would use again the reconstructive nature of the memory, but

[1] Dept. of Software and Computing Science Education, Charles University in Prague, Ke Karlovu 3, 12118, Prague, Czech Republic.
[2] Institute of Psychology. Academy of Sciences, Politických vězňů 7, 11000, Prague, Czech Republic

now going from the reason why you went there (tooth ache, long-distance travel) and asking when this cause occurred. Finally, you may recall that it was approximately in May last year, event though the right answer was the 5th June.

Data from various experiments (e.g. [15]) show that the quality of recall can differ on different time scales (hours, days, months). Moreover, the data show that when asked to date an event people usually use various socially agreed general time patterns when speaking about the time ("morning", "evening", etc.) instead of from 5:40 to 6:38.

These points can help us with the design of a timing component for an episodic memory for virtual characters. People do not recall time of events exactly (e.g. [14, 16]) which suggests that an IVA does not need to store time information in the exact manner either. The comprehension via time patterns also contributes to parsimonious design of an IVA's long term memory system since the memory resources can be saved by storing only the required time information. Moreover, the stored time information can mimic better socially agreed time patterns which are used by people in everyday communication.

We have already developed a timing component for an episodic memory for IVAs [17] and integrated it with the rest of our previous episodic memory model [18], producing a novel connectionist model of episodic memory for virtual characters. The important feature of the model is that its timing component is able to learn the time patterns, such as "early afternoon", automatically based on the IVA's life style. The time patterns the model learns cluster contextual information with a period of time; for instance, the context of some morning activities with time from 7.00 to 8.30 a.m. Those time patterns are anonymous – they are not labelled as "morning" or "afternoon" directly. The labelling is produced when the IVA uses a pattern as a part of an answer; then the name of the most matching socially agreed time pattern is assigned. For instance, the abovementioned previous cluster would be labelled as "morning".

In sum, our objective has been to enhance IVAs with the following skills:

1. Learning of internal representation of time patterns.
2. Ability to comprehend human-like time concepts.
3. Ability to retell stories using human-like timing.

However, there were two problems with the model we developed. First, the key parameters – the forgetting rates – remained unspecified. Second, it was not clear whether the time patterns the model learns were believable. The time patterns are used both in the query construction and in the recall. Thus there are two areas of believability to cover. In this paper, we address the former one: the query construction. The issue here is that we did not know which types of time-cued questions are typically asked by users of IVA applications, what time patterns they use and what they are concerned about. Thus, we conducted a questionnaire study to shed light on this issue. We tried to determine which type of time patterns are given by users to IVAs, e.g. what time information is usually carried by a query. The results were incorporated in the internal memory organization allowing IVAs to use these time patterns as the input for the recall. As a side product, the incorporation of the time patterns allowed for easier use of such patterns when an IVA was retelling a story about its past.

The forgetting rates remain as a work-in-progress. There is no consensus in psychology which forgetting curve describes best behaviour of human memory [19]; nevertheless, there are many psychological works which can help IVA developers to start to address this issue (e.g. [16, 19, 20, 21, 22]). A future work is a thorough believability study on human subjects in which IVAs telling stories about themselves based on their episodic memory would be subject to a kind of Turing test.

The main goal of this paper is to present the findings from the questionnaire study and to review the memory model. The results suggest that people do prefer vaguer time patterns over exact time specification (with some exceptions) and that it is possible to design and implement a timing component compatible with those time patterns.

The paper is organized as follows. Chapter 2 summarizes methods and results of the questionnaire experiments concerning time-cued questions and time patterns. Chapter 3 reviews the prototype of the timing component of the long-term episodic memory followed up by results obtained from several experiments. Additional materials are available on-line on our website[3] and in [17].

2 QUESTIONNAIRE

The questionnaire was carried out in three consecutive rounds. We started with asking several people to write down a set of questions which they would ask an imaginary NPC (a non-player character) in a role-playing game. This preliminary phase had 7 participants, all men, undergraduate students, second to fourth grade. From final list we picked questions containing some time-related cues and recalls of past events. Based on these questions, we have created a set of questions for Experiment I. The results of Experiment I were used to refine the scenario and the set of questions for Experiment II.

2.1 Experiment I

Questionnaire and Hypothesis

The questions from the preliminary phase were extended by a set of similar questions with altered specification of time (e.g. *"What did you do yesterday afternoon?"* → *"What did you do on Friday at 13:15?"*). The aim of this alternation was to model various types of time-cues. Supposedly, the questions with vaguer time information would be favoured by the participants. When designing the questionnaire we wanted to cover a larger scope of possible tasks and time precision preferences. Therefore we introduced the participants to a rather general imaginary scenario rather than a more developed specific one. The latter would be also a good option, but it would be necessary to add a lot of additional information to the instruction to describe the setting and related tasks.

Method

The questionnaire was distributed to 30 volunteers (27 men, 3 women). The RPG game experience varied across the sample

[3] http://artemis.ms.mff.cuni.cz/pogamut/tiki-index.php?page=Episodic+memory+for+virtual+agent

from little/none to large. They were mostly undergraduate students of computer science. They were given approximately 10 minutes for the task.

The participants were introduced to the situation by a short description of the model scenario, which was roughly as follows: *You are a player of an RPG game which takes place in a vast virtual world. There you can move freely and communicate with its inhabitants. Those inhabitants are controlled by the computer and they are able to perceive the surrounding world, reason about it and answer a variety of questions. You have just arrived to a new town and you are going to acquire a quest. To this end, you first talk to a merchant.*

Then the subjects were given a list of questions. The task was to rate these questions on the scale from 1 to 5 where 1 indicated a weird question while 5 indicated a valid question one would ask a virtual character. The exact version of the questionnaire with the task description can be found in Appendix.

Results

Figure 1 depicts the average response with standard deviation for each question. The evidence that would support our hypothesis, that exactly timed questions would score lower than vaguer questions, was not strong enough. Nevertheless, we were able to distinguish *post hoc* several categories among the questions which can explain some unpredicted ratings. The scenario concerned the acquisition of a quest, which can explain high values for questions related with criminality, orcs, guards and healing potions (six questions – see Appendix –, mean = 4.14, SD = 0.95). But we have also found that most of the artificially added questions with the accurate time information were rejected (four questions, mean = 1.79, SD = 0.96) as well as closely personal questions (three questions, mean = 2.23, SD = 1.18) as they were not important to the interrogator. On the other hand, personal questions with some relation to a possible quest or the overall situation were rated high (three questions, mean = 3.91, SD = 0.96).

We have not found any notable relation between the amount of RPG game experience of participants and the ratings.

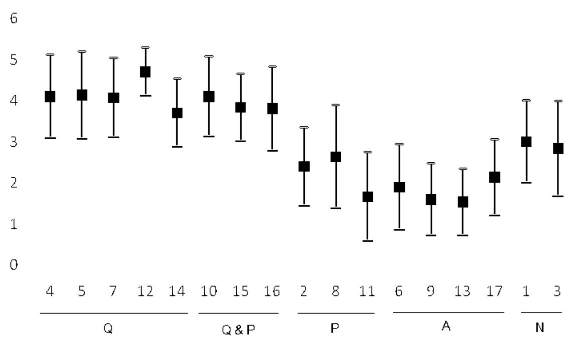

Figure 1. Mean question ratings. Whiskers indicate standard deviation from the average in the middle. Y axis represents rating (1 indicates a weird question, 5 a valid question). X axis contains numbers of questions as they appeared in the questionnaire. First 5 questions are quest related (Q), next 3 are personal and quest related (Q & P). Then follow 3 personal questions (P), 4 artificial questions (A) and 2 neutral (N).

Discussion

Experiment I showed that we are pursuing the right direction. Nevertheless, its results were strongly influenced by the proposed scenario situated in an RPG game, which usually includes goblin raids, magic potions etc. Thus we have conducted another experiment with modified settings as well as with the altered set of questions.

2.2 Experiment II

Questionnaire and Hypothesis

The previous results allow us to form the following hypothesis. We divide the time questions into three categories based on the specificity of the time information:

1. Exact time is given or demanded & thorough recollection is demanded: an answer is supposed to contain exact time information specified in numbers, e.g. 13th April morning or at 13:17; a question is supposed to contain such time cues ("What were you doing the 13th April morning?")
2. Neither category 1 nor 3: e.g. questions targeting events more than one week old *plus* demanding their dating on the scale of parts of a day (the morning of last Christmas), etc.
3. Vague time cue is given & recollection is easy: i) questions that contain as a time cue a recent "past day" time concept such as "yesterday" or a combination of a "multi hour" fuzzy time concept (MHTC) plus a recent "past day" time concept such as "yesterday morning"; ii) a question on a repetitive activity that contains as a time cue a particular day of week or a particular day of week plus MHTC (e.g. Monday evenings); iii) a question that contains a vague "multi day" time cue without specifying more detail such as "last week", "recently", "in last few days"; or iv) a when-question that can be answered easily and using vague time concepts ("When were you on your last holidays?" – "Last winter.")

The assignment of questions into the categories is necessarily subjective. Three of the authors of this paper (CB, OB, JL) voted on each question (i.e. they selected a category number or "I don't know") and only questions on which at least two of them agreed upon the category and the last voted "I don't know" or his preferred category for this question differed from the other votes at most by 1, was assigned into that category. Several questions were removed from the questionnaire in the design phase based on the result of this voting procedure. We did not use variants of the same question adapted for categories 1 or 3 in order to conceal the hypothesis to the participants.

We presume that people would prefer questions from category 3 over questions from category 2 over questions from category 1.

Method

The questionnaire was distributed by email to former participants of Summer School of Mathematics and Physics which is a camp for talented young interested in science.

Another group of participants was formed by undergraduate students of computer science and psychology and the last group consisted of several graduates. The response ratio was about 30%. We have gathered 24 responses (7 women and 17 men ranging from mostly 19 to 28 years, with 3 older). All subjects were distinct from the subjects of Experiment I.

This time we wanted to avoid tight coupling with a concrete situation and an environment. Thus we have described the world more vaguely and proposed several scenarios to stimulate imagination of participants:

There are several towns in the virtual world. The situation can fall, for instance, into one of those scenarios:

1. *You live in one of the towns and you go regularly to the pub. You engage in a conversation with a newcomer trying to get some news and personal information about him.*
2. *You have come to a new location and figure out how things go here, what are they up to here.*
3. *You have met another player. You grab a virtual coffee and chat a little.*
4. *You chat with a merchant in his boutique.*

The description was followed up by the set of 36 questions and the participants were asked to rank them on the same scale from 1 to 5 as in the previous experiment. There were 11 questions of first category, 9 of the second, 13 of the third and three other questions with no time information (i.e. lures).

Results

Figure 2 depicts the mean and standard deviation of ratings for the three categories of questions. One-way ANOVA test was run on the data and confirmed that the means of all categories are non-equal ($p < 0.001$). There is clear evidence that exactly defined questions are less convenient than vaguely specified questions. Nevertheless, we have discovered few surprising discrepancies. First, one question from the 1st category "*Do you know when we exactly went to sleep after that marathon run last Saturday?"* (mean = 3.1, SD = 0.95) was rated relatively highly, possibly thanks to the compelling context. Second, the personal questions were again rated lowly in comparison with other questions in the same category despite their vague timing, for instance, two questions from the 2nd category. "*How long did you eat the lunch last Thursday?*" (mean = 1.3, SD = 0.46) and the question "*What did you do last Friday after breakfast?*" (mean = 2.1, SD = 0.97). It is possibly the unimportant context that decreases the rating of the questions. Third, one question from the 2nd category scored over 4 on average, despite the given time information was: *"Monday two weeks ago"*. This may indicate that the vagueness of time information can include even such patterns.

The questions in the third category were rated consistently high.

We again asked participants about their experience with virtual worlds; however, we did not find notable dependence between the ratings and the experience level of respondents. The only notable discrepancy was the "marathon" question, which received good ratings from non-players (mean = 3.6) whilst obtaining relatively low marks from "gamers" (mean = 2.6). See Appendix for complete overview of results.

As was mentioned above the assignment of questions into categories was a subjective process. Differences between categories can be best illustrated on some border cases. Consider question no. 4: *"What were you doing Wednesday the previous week?"*, that was assigned to category 2. If it had addressed Wednesday from this week, it would have fallen to category 3. On the other hand, question no. 7: *"Are you doing something on Sunday between 8 p.m. and midnight? We can go for a beer."* was assigned to category 2 even though it contained exact time information, because the context made it clear that we were actually asking for the evening. A similar type is question no. 35: *"Are you doing something Tuesdays between 17:15 and 18:45? We can play squash."* This question was also assigned to the second group despite the exact time information.

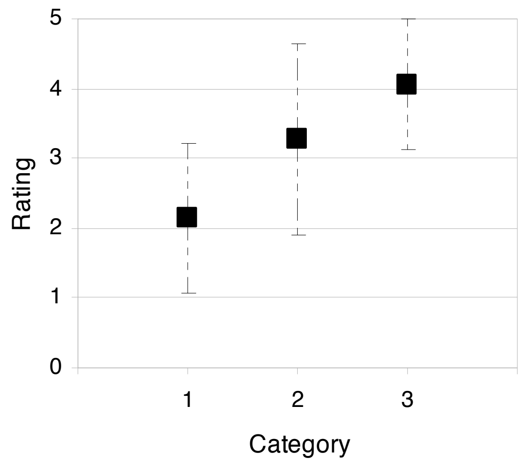

Figure 2. Mean category ratings. Whiskers indicate standard deviation, 1 indicates a weird question, 5 a valid question.

Discussion

The collected data is consistent with the proposed hypothesis that people prefer questions using only vague time specifications. The questions with exact time were rated lower, with an exception of "Squash question", where the exact time is probably of high importance in everyday life and the "marathon run" question where the context probably influenced its rating. The vaguely specified questions from the 3rd category were rated highly.

Ultimately, the gathered data leads us to the conclusion that people would rather ask IVAs using vague time cues than exact ones. Now we know that IVAs should be capable of processing vaguely specified questions. Therefore they need more profound internal understanding of human-like time cues to be able to respond believably. Additionally, we can draw following points from the experiments. First, an IVA typically does not need to store its personal information in a great detail as the user's interest lies usually in other domains. Second, the timing that belongs onto the scale of parts of a day ("morning", "evening") seems to be sufficient for most cases.

Proceedings of the Remembering Who We Are – Human Memory for Artificial Agents Symposium, Mei Yii Lim and Wan Ching Ho (Eds.), at the AISB 2010 convention, 29 March – 1 April 2010, De Montfort University, Leicester, UK

For future research, it would be interesting to investigate in detail ordinary human answers to these questions, the influence of context and relevance of questions to the described scenario, and how specific or vague place/object references are accepted in human dialogues. Answers to these issues can help with further constriction of requirements on episodic memory models for IVAs.

3 CONNECTIONIST EPISODIC MEMORY DESIGN

The purpose of this section is to review the timing component for the episodic memory for IVAs we have developed. For the space constraints, we cannot introduce all detail; these can be found in [17]. We will first review our agent architecture, second, the general episodic memory system, and third, the timing component. The chapter concludes by the summary of experiments made on the model.

3.1 Agent architecture

Our IVA's overall architecture is depicted in Figure 3. It is a reminiscence of a classical cognitive AI architecture, by which many virtual agents have been inspired. Our IVA is driven by *hierarchical reactive planning* with behaviour represented by *AND-OR trees*. 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 can be achieved by adopting some sub-goals. The IVA needs to perform only one task to achieve a goal, provided there is no failure (hence, OR nodes); but to fulfil all sub-goals to solve a task (hence, AND nodes). The tasks that cannot be further decomposed are *atomic actions*, i.e. action primitives. Every task may need several resources to be performed, i.e. *objects*. Every top-level goal has its *activity level* based on drives, external events, and a schedule. The competition among the goals based on this level 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 *task field* of the *visual short term memory*. One goal can interrupt another, in which case the tasks of the interrupted top-level goal are remembered and can be resumed after the more important goal is achieved. From the AI standpoint, this mechanism capitalises on the BDI framework [23].

The visual short-term memory holds templates of objects seen that passed through a simple threshold-based *attention filter*. Every object is regarded as a tool for action, i.e. it *is* a set of "affordances" [24], meaning it possesses pointers to the tasks it can be used as a resource for. These pointers are perceived directly by the IVA when observing his environment. Objects in experiments we have been running are state-less for the sake of simplification, though our simulations allow the objects to have states as well. The *memory field* can also temporarily hold information about an object recalled from the long-term memory.

This architecture has several long-term memory appendages, such as spatial memory (LTSM) and episodic memory (LTEM). Here, we are concerned only with the episodic memory.

3.2 Long-term episodic memory

Recall that our present episodic memory model with the timing component extends the first generation of the model presented in [18]. While the old model was symbolic, the present one is connectionist. The present model is composed of a three-lobed neural network (Figure 4).

The core lobe is the *long-term episodic memory* (LTEM), which stores past episodes of the IVA's life, the flow of events. The memory contains a forest of AND-OR trees representing all possible tasks the IVA can perform (Figure 4 – middle). This structure mirrors IVA's behavioural representation. Note that this was the only part of the previous model [18]. In present version of the model, this core lobe is surrounded by other lobes responsible for timing.

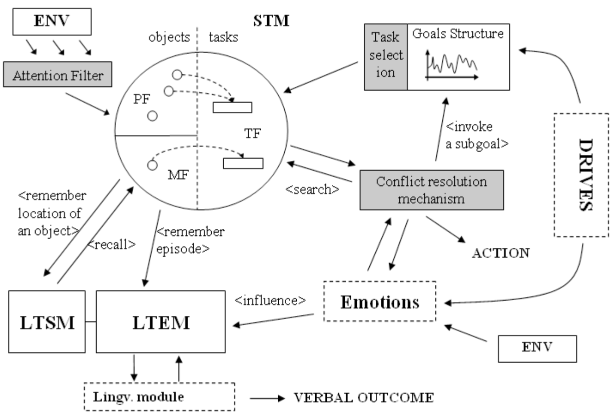

Figure 3. Outline of our IVA design. He receives inputs from the *environment* (ENV) which fills up the *perception field* (PF) of the *short-term memory* (STM). The decisions are made according to the activation of *goals* and *drives*. Tasks in progress are represented in the *task field* (TF). *Memory field* (MF) stores locations of resources obtained from the *long-term spatial memory* (LTSM) on request. Personal episodes are stored in the *long-term episodic memory* (LTEM). Forgetting in the LTEM is influenced by a simple valence-based *emotion model*.

The first surrounding lobe is a two-layered neural network. The first, input layer consists of *context nodes*, which describe the agent's internal or external state (i.e., context), and *Cartesian nodes*, which represent objective time (i.e., their activation changes regularly over time). Cartesian nodes can be also conceived as biorhythms. This layer is connected with the layer of daytime concepts. *Concept nodes* associate together particular context the IVA is situated in with some period of daytime, thus forming new time patterns. For example, if usual afternoon activities of our IVA comprise of studying for exams, some concept nodes will be linked with context nodes for studying and Cartesian nodes for time between 14.00 and 17.00 hrs. In the end, concept nodes should represent the desired vaguer human-like time patterns. Importantly, unlike the hard-coded layer of Cartesian and context nodes, the concept nodes are learned automatically during the course of simulation using Hebbian learning with subtractive normalization (see [17] for details).

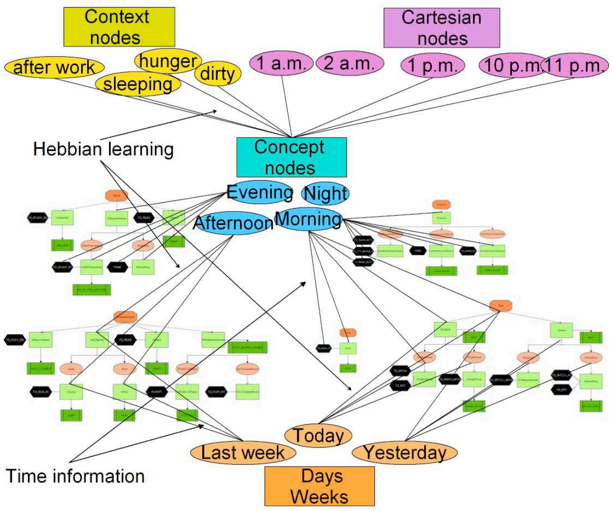

Figure 4. Episodic memory design outline. The context layer (top left) and Cartesian layer (top right) are connected with concept nodes (centre). Concept nodes are connected with AND-OR trees of goals/desires (middle bottom). These are also connected with the day/week nodes (at the bottom).

The second lobe contains *day nodes* (bottom at Fig. 4). These represent past days, such as "yesterday", or abstractions for multiple-days and parts of weeks (such as "first two days of the week" or "last day of workdays" etc.). Note, that some psychological data suggest that humans use these abstractions [15]. These nodes are hard-coded.

All Cartesian, context, concept, and day nodes work as rate-coded neurons with a <0, *m*> response function, where *m* is a tuning constant, which is typically from <2, 4> (see [17]).

The AND-OR trees are responsible for the actual episodic entry storage. During storage, the nodes of the AND-OR trees are interconnected with two kinds of nodes. First, with concept and day nodes, representing the time information, which can be later used during recall. Second, with nodes representing objects used in pursuing tasks and scenes where the tasks took place. The storage, i.e. building of links, happens continuously along the course of IVA's actions. The links are weighted and the weights are built again on the Hebbian basis (see [17] for details).

Note that events the core lobe stores have various grain size, which allows for *gradual forgetting*: unimportant details of episodes can be "eaten away" from the bottom of the AND-OR trees. In the prototype implementation, there is a given threshold for each day which is decreasing over time. Every night, the network representing "today" is labelled as "yesterday" and analogically for the other days. Then the weights are decreased by an arbitrary coefficient. Weights which are weaker than an arbitrarily chosen threshold are discarded and their energy (weight) is distributed amongst remaining weights which promotes storing of important (stronger) connections. For example, one can originally remember that he was cooking a goulash yesterday morning, including all subtasks, but later forget the subtasks, keeping only the high-level information about cooking. To our knowledge, present episodic memory systems for IVAs do not have this feature. Moreover, it allows the IVA, to some extent, to blend similar episodes into one (see Sec. 3.4). The details of the forgetting are described in [18] and other concept related to our memory systems in [5, 25].

3.3 Labelling of concept nodes and recall

So far, we have discussed only the structure of the episodic memory system, development of time concepts and storage. This section will cover the use of the built network for episodic recall. For recall, the memory structure is not conceived as a regular neural network, but as a spreading activation network. This means that we can inject activity into a node and propagate this activity to neighbouring nodes. The activity is propagated proportionally to the weight of each link. In the present version of the model, the activity is propagated only to direct neighbours (for the sake of analysing the experimental results).

Let us imagine that the IVA has already learned concept nodes (this takes approximately 10 days (of the time in the virtual world) with the current setting of parameters) and there is already stored some information in the network. Now, the user asks a question. Every question contains some cues – time, place, activity. Those cues are conceived as inputs for our memory system: basically, nodes corresponding to these cues are activated. The activity is propagated through the network from neurons representing question cues and the most active neurons of the demanded category are returned. For instance, question: *"What were you doing in the evening two days ago?"* infuses activity to the neurons for *evening* and *two days ago* and returns the most active activities; in example depicted on Figure 5 these are *swimming* and *dinner*.

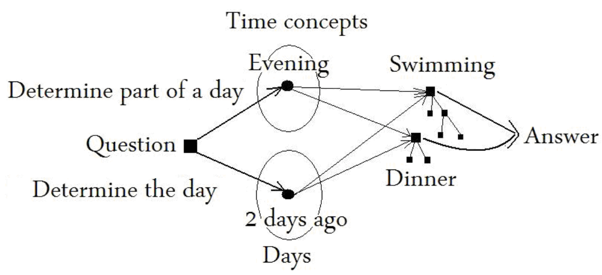

Figure 5. A query example. The arrows depict propagation of the activity. The activity is infused into the neurons for *evening* and *two days ago* and the most relevant answer based on the activity level is "*swimming* and *dinner*".

With respect to questions on time information, there is one issue. As we have mentioned in the previous section, concept nodes are anonymous, internal representation of time concept clusters in the IVA's mind. The recall returns which concept nodes are relevant to the question but it cannot provide their names. Therefore, we need a mechanism which would translate those internal time concepts into socially agreed time patterns. To this end, we use the links between Cartesian nodes and concept nodes. First of all, we define a set of six parts of a day – labels (morning, noon, afternoon, evening, late evening and night). Each label confines a fixed time interval. Then we measure the correspondence of concept nodes to one of those labels: the most appropriate label is assigned to each of the concept node. The measure of correspondence is obtained from weighted links between Cartesian nodes and concept nodes.

First, Cartesian nodes are assigned to human categories. For example, Cartesian nodes for time from 6 a.m. to 9 a.m. fall to "morning". Then we activate all Cartesian nodes from a particular category and propagate their activity to the concept nodes, measuring the concept nodes' activation. Each concept node is assigned the category that activated the node mostly. A crude natural analogy to this process is a child being taught by her parents to name various parts of day.

A notable feature of our mechanism is that the recall can be error prone. For instance, tasks and objects linked with the "after breakfast" node for "yesterday" can be merged with tasks and used objects that were linked with the "after breakfast" node the day "before yesterday". Thus two episodes can be blended during the recall (when the recall is cued by time information).

3.4 Experiments

The design of the memory system was followed up by a successful implementation of the prototype [17] on the platform for fast prototyping of virtual agents Pogamut 2[4] [26]. The implementation was tested by a set of experiments. The experiments included tests of memory accuracy, memory space demands, episode blending and the study of impact of various settings of context and Cartesian nodes on the memory performance. The simulation usually lasted, from the standpoint of the IVA, two to four weeks. The IVA lived according to several different lifestyles (student, millionaire, travel salesman).

Lifestyles usually contained few fixed, routine activities and few dynamic, occasional activities. The world of the IVA comprised of 6 rooms and 27 different objects and places and was situated in the 3D FPS game Unreal Tournament 2004. The agent had 10 different activities to perform. The student's plan was filled with lectures during the week days mornings followed by various sport activities in the afternoons and cultural events in the evenings. The millionaire had a less structured life filled with entertaining activities. The travelling salesman worked mostly all the time with only occasional divergence in the plan. Each character had different time frame of the day (alarm clock settings).

The results we gained so far suggest that:

1. The network is able to learn and use the concept nodes properly and these concepts are qualitatively similar to those used by humans [15].
2. The memory is error-prone and blending can occur under some circumstances, e.g. an IVA who was watering a garden every evening cannot recall details of any particular watering episode when cued by time, but it can recall both that it was watering evenings last week as well as some detail of each watering when cued with time plus other details of the episode. However, this behaviour remains to be investigated on a rigour basis and the whole mechanism should be tuned appropriately; a future work.
3. The forgetting mechanism can ensure smooth forgetting of details and unimportant episodes.
4. The activation of Cartesian nodes should overlap to create a more compact time clusters represented by concept nodes.
5. Context nodes play an important role in the learning process. They support the diversification of the Cartesian nodes; more concept nodes are formed for periods in which activities change often (i.e. with different context nodes being activated) comparing to periods without changes of activities (e.g. a night). This is a desirable outcome.

To conclude, the conducted experiments revealed some key parameters of the learning process (i.e. the setup of Cartesian and context nodes) and interesting dependencies (i.e. the importance of context nodes). They also helped to verify the design of the memory system as they showed it is capable of recall on time-cued questions. The experiments also demonstrated the capacity of the memory to blend episodes and to forget details while keeping the important parts of memories.

4. CONCLUSION

The goal of this paper was twofold. First, the aim was to identify requirements for time representation underpinning episodic memory systems for IVAs, enabling conversations featuring believable time patterns. Second, the aim was to review our memory model adhering to theses requirements.

The data gathered from the two questionnaire rounds suggest that people would use rather socially established time patterns than exact time or date when questioning virtual characters. The time patterns can be parts of a day such as "morning", or denotations of larger time intervals such as "last week". Additionally, the data suggest that human users would not ask virtual characters for detailed depiction of personal concerns, unless they are directly relevant to the users' interest.

We have designed, implemented and tested a novel model of episodic memory enhanced by a timing component. The results of experiments with this model suggest that the model enhance IVAs with ability to understand and work with human-like time patterns. Its connectionist basis allows IVAs to mimic some phenomena of human memory like episode blending, error proneness and gradual forgetting of details. In future, this episodic memory module can be used in many role-playing games and other applications featuring virtual characters inhabiting their worlds over larger periods of time.

APPENDIX

Appendix is available online on address: http://artemis.ms.mff.cuni.cz/main/papers/Timing-Appendix.pdf.

ACKNOWLEDGMENTS

This work was partially supported by the research project MSM0021620838 of the Ministry of Education of the Czech Republic, and by the student research grant GA UK 21809 and by the grant 201/09/H057.

REFERENCES

[4] http://pogamut.cuni.cz/

[1]E. Tulving, W. Donaldson. *Organization of memory*. New York: Academic Press. (1972).

[2]A. Baddeley, M. J. Conway. Aggelton. *Episodic Memory: New directions in research,* Oxford Uni. Press. (2001).

[3]M. Wooldridge. *Reasoning about rational agents*, MIT press. (2000).

[4]B. Loyall. Believable Agents: Building Interactive Personalities. Tech report CMU-CS-97-123. Ph.D. dissertation, Carnegie Mellon University. (1997).

[5]C. Brom, J. Lukavský. Towards virstual characters with episodic memory II: Episodic memory strikes back. In: *Proc. Empathic Agents, AAMAS workshop*. 1 - 9. (2009).

[6]W. C. Ho, S. Watson. Autobiographic knowledge for believable virtual characters. In: *Proc. of Intelligent Virtual Agents*, LNCS 4133. 383-394. Berlin, Springer-Verlag. (2006).

[7]G. Castellano, R. Aylett, K. Dautenhahn, A. Paiva, P. W. McOwan, S. Ho, Long-term affect sensitive and socially interactive companions. In: *4th Int. Workshop on Human-Computer Conversation*. (2008).

[8]I. Doherty, C. O'Riordan. Toward More Humanlike NPCs for First-/Third-Person Shooter Games. In*: AI Game Programming Wisdom IV*. Charles River Media. 499 – 512. (2008).

[9]A. Nuxoll. Enhancing Intelligent Agents with Episodic Memory. PhD thesis, The University of Michigan. (2007).

[10]W. Dodd. The design of procedural, semantic, and episodic memory systems for a cognitive robot. Master thesis. Vanderbilt University, Nashville, Tennessee. (2005).

[11]W. Ho, K. Dautenhahn, C. Nehaniv. Computational Memory Architectures for Autobiographic Agents Interacting in a Complex Virtual Environment. In: *Connection Science*. (2008).

[12]J. Rickel, W. L. Johnson. Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control. *App. Artificial Intelligence* 13 (4-5): 343-382. (1999).

[13]J. Dias, W. C. Ho, T. Vogt, N. Beeckman, A. Paiva, E. Andre. I Know What I Did Last Summer: Autobiographic Memory in Synthetic Characters. In: *Proc. of ACII*. 606–617. Berlin, Springer-Verlag. (2007).

[14]W. J. Friedman. Memory for the Time of Past Events. In: *Psychol Bull* 113(1): 44 – 66. (1993).

[15]S. F. Larsen, C. P. Thompson, T. Hansen. Time in Autobiographical Memory, In: *Remembering Our Past: Studies in Autobiographical Memory*, New York: Cambridge University Press, 129-156. (1995).

[16]W. A. Wagenaar. My memory: A study of autobiographical memory over six years. In: Cognitive Psychology. Vol 18(2), 225-252. (1986).

[17]O. Burkert. Connectionist Model of Episodic Memory for Virtual Humans. Master thesis. Charles University in Prague. (2009).

[18]C. Brom, K. Pešková, J. Lukavský. What does your actor remember? Towards characters with a full episodic memory. In: *Proc. ICVS*, LNCS 4871, Springer-Verlag. (2007).

[19]R. A. Chechile. Memory hazard functions: A vehicle for theory development and test. In: *Psychological Review*. Vol 113(1), 31-56. (2006).

[20]P. H. Bahrick, E. Phelphs: Retention of Spanish vocabulary over 8 years. In: *Journal of Experimental Psychology: Learning, Memory, and Cognition*. Vol 13(2), 344-349. (1987).

[21]J. R. Anderson, R. Milson. Human memory: An adaptive perspective. In: *Psychological Review*. Vol 96(4), 703-719. (1989).

[22]O. Furman, N. Dorfman, U. Hasson, L. Davachi, Y. Dudai. They saw a movie: Long-term memory for an extended audiovisual narrative. In: *Learning & Memory 14*, 457–467. (2007).

[23]M. E. Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press: Cambridge, MA. (1987).

[24]J. J. Gibson. The Theory of Affordances. In: *Perceiving, Acting, and Knowing*, Eds. Robert Shaw and John Bransford. (1977).

[25]Brom, J. Lukavský. Episodic Memory for Human-like Agents and Human-like Agents for Episodic Memory. In: *Technical Reports FS-08-04, Papers from the AAAI Fall Symposium Biologically Inspired Cognitive Architectures*. Westin Arlington, VA, USA, pp. 42-47, AAAI Press, (2008).

[26]O. Burkert, R. Kadlec, J. Gemrot, M. Bída, J. Havlíček, M. Dörfler, C. Brom. Towards fast prototyping of IVAs behavior: Pogamut 2 In: *Proceedings of 7th International Conference on Inteligent Virtual Humans*, LNCS Vol. 4722. Paris, France. Springer-Verlag, Berlin. (2007).