Timing in Episodic Memory for Virtual Characters

Cyril Brom, Ondřej Burkert, and Rudolf Kadlec

Abstract—**Recently several episodic memory models have been developed for virtual characters to increase their believability. However, none of these models addresses the issue of plausible timing of events. Here we present a model that addresses this issue. We introduce a prototype implementation and discuss the psychological underpinnings. Then we demonstrate that the model is able to mimic some psychological phenomena such as blending similar episodes.**

I. Introduction

Enhancing cognitive abilities of non-player characters (NPCs) can increase game-play quality in many videogames [1]. Episodic memory is one such ability. In psychological terms, *episodic memory* (EM) [2, 17] is an umbrella term for memory systems operating with representations of an entity's personal history. EM traces are related to particular places and moments and connected to subjective feelings and current goals. The EM is 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 (e.g. "France is a part of Europe"). The latter covers processes related to skill learning.

EM modelling has gone almost unstudied in the context of videogames. However, because there is now a growing interest in this issue in the neighbouring field, the study of intelligent virtual agents (IVAs), it may be soon possible to develop NPCs with EM abilities capitalising on the knowledge gained by the IVA research. Examples of skills an NPC (or an IVA) can possess that demand some facet of EM include, but are not limited to: a) general giving of information to a user based on the past history of the NPC's interaction with the virtual world, b) imitating a dialog between two NPCs, again, based on their personal history, c) remembering a course of interaction with a user, notably remembering the course of a dialog, d) learning [3]. These skills would increase the believability of many NPCs in shooter games and role-playing games (RPGs). For instance, NPCs could answer questions like "what happened in your shop yesterday?" based on their history and not using pre-scripted answers.

This work was partially supported by the student grant GA UK 21809, 44910 and by the research project MSM0021620838 of the Ministry of Education of the Czech Republic.

Cyril Brom is an assistant professor at the Faculty of Mathematics and Physics, Charles University in Prague, Malostranské nám. 25, 118 00, Praha 1, CZ (phone: +420221914217; e-mail: brom@ksvi.mff.cuni.cz).

Ondřej Burkert is a former student of the Faculty of Mathematics and Physics, Charles University in Prague.

Rudolf Kadlec is a PhD student there.

EM modelling is a new topic even in the context of IVAs. Research so far has focused mainly on developing proof-of-concept implementations [4 – 9]. Only some of the models were included in final applications [4, 5, 9]. Evaluations of multiple models against each other are rare [see 6, 10 for the first swallows], as well as evaluations of models against psychological data and/or during a believability study in which subjects are asked whether an IVA with EM abilities appears human [see 11 for an exception]. Addressing technical issues like consumption of computational resources is also rare [see 7, 8].

Additionally, because human EM is a multi-faceted phenomenon, many important aspects of EM have not been modelled yet. One such missing trait is the ability of estimating the time when an event happened. Without the ability of time events, human life would be dull. For instance, imagine yourself recalling that you brushed your teeth in the morning, but not whether this happened a minute ago or three days ago. IVAs should be equipped with a similar ability for many applications. Obviously, one can store precise time information ("I was cooking yesterday from 1:12 p.m. to 2:37 p.m."). However, this is neither psychologically plausible, nor would it be believable to the audience. For instance, time information in human memory often deteriorates as episodes get older.

The goal of this paper is to present a computational module for believable timing by IVAs (*timing module* throughout). The model is primarily intended for NPCs from role-playing games, though it can be used in any simulation featuring long-living IVAs. The model has three notable features: 1) It stores information using socially agreed time patterns (e.g. "morning", "after lunch"); 2) These time patterns are learnt, to some extent, automatically based on the IVA's history, using an artificial neural network with Hebbian learning; 3) The module is able to blend similar episodes during forgetting ("I remember I was gardening some evenings last week, but I don't know which days exactly. I also remember some perceptual details, but I don't know what detail happened on which day."). The present version of the model works with time periods of weeks.

The paper proceeds as follows. Sec. II reviews what is known about human timing and presents general requirements on the timing module for IVAs. Due to the novelty of the topic, we deem it appropriate to include such a section here. Sec. III introduces architecture of our agent. Sec. IV presents the memory model and its timing module. Sec. V presents the prototype implementation and details a part of the evaluation. The evaluation is a compromise between what is possible and what would be optimal. First, there is no model against which we could compare our model. Second, perhaps surprisingly for some, up to date psychological findings are not sufficient to provide information needed for quantitative evaluation of the model against human data [3]. Third, to conduct a rigorous believability study, it would be necessary to develop an appropriate infrastructure, e.g. a dialog system and graphical content. Thus, in our evaluation, we focused on demonstrating that the timing module is able to mimic some phenomena of human timing qualitatively. For brevity, we detail here only the test that concerns blending similar episodes. Other experiments are described in [12]. Additionally, results of a precursor to a believability study concerning whether RPG players prefer to use time patterns to NPCs when speaking can be found in a complementary paper [13].

II. Drawing s of Timing

In general psychology and neuropsychology offers useful metaphors, conceptual models and human data that can help IVA (and NPC) designers to answer many question posed by EM modelling. Notable examples include clarifying the notions of *short-term* vs. *long-term* memory, specifying what should be remembered by an IVA and why, or specifying the processes of *storage*, *retrieval*, and *maintenance,* which includes *forgetting*. However, beyond conceptual models and metaphors, the issues become hazy. The modellers must basically come up with many *ad hoc* solutions to fill the gaps and specify their models in precise detail [3]. This will be exemplified in this paper on the phenomenon of timing.

A comprehensive review of psychological theories of human timing is given in [14, 15]. These reviews make it clear that human timing is not perfectly resolved by psychology yet. Still, there are several psychological theories on human timing as well as evidence from experiments that is of importance to IVAs' designers.

Perhaps the simplest theories are based on the idea that the time of events, e.g. date, is coded explicitly: these are called *time-tag* theories [14, 15]. Strengths and accuracy of these tags possibly decay with time. However, humans are poor at giving the exact time of events as well as using time information as a cue for retrieving events [16]. Data suggest that the time of only very important events is learnt, so-called *landmarks* [14, 15], and these events are later used as reference points. A different family of theories capitalises on the idea that events are chronologically ordered via "pointers", so-called *event-succession* theories [15]. In this view events need not be time stamped. Yet other kinds of theories emphasise the *reconstructive* nature of EM: reconstructive recall combines general semantic knowledge, knowledge of one's lifestyle and personal episodes to approximate the time of an event.

The outcome of this debate [14, 15] is that humans use multiple techniques for dating past events, among which the reconstructive way is the most important one.

For instance, if you try to recall what you did on "December 20 last year", you would automatically convert the date to "five days before last Christmas" and use Christmas as an important time landmark. You would remember that you were working two days before that Christmas, resulting in the refinement of the time information to "Saturday before that Christmas". Thus you would conclude that you were probably shopping for gifts. Another example: assume you are recalling when you visited your dentist the last time. It is likely that you would again reconstruct the information, but 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, even though the right answer is June 5. Yet another example: you are asked when the event happened in which your brother broke a window. Because you remember that the window was broken with a snowball, you would conclude that it was during winter.

Additionally, data from various experiments show that the quality of recall can differ independently on different time scales (hours, days, months) [15]. For example, it was found in an experiment on the dating of events that subjects made errors that tended to be in multiples of 7 days, suggesting that a day of the week was represented more often than an exact date [15]. Data also show that when asked to date an event, people usually use various socially-agreed time patterns ("morning", "evening", etc.) instead of exact information ("from 5:40 to 6:38"). Thus, the reconstructive nature of recall is supplemented by *temporal patterns* or *schemata*: "...the evidence is overwhelming that the primary process is reconstruction based on schemata for time" [15, p. 130]. Our own data support this finding [13]. We conducted a simple study on human subjects using questionnaires (age=19-28; n=24; 17 males) asking whether the subjects would prefer to ask NPCs questions regarding their past using time patterns ("morning") or exact timing information ("6:38"); there was a strong tendency among subjects to prefer the former.

The tentative conclusions we draw from our review of psychological literature are:

1. The timing module for IVAs should link consecutive episodes according to the event-succession theories. These links can decay until most of them are forgotten.
2. The timing module for IVAs should use *time concepts*; i.e. socially-agreed time patterns like "yesterday" or "night". These should be used not only when speaking about events, but also for the representation of time. Time concepts are analogies of temporal schemata used in [15].
3. The timing module should exploit time concepts that are overlapping ("early morning" vs. late morning") and hierarchically nested ("morning" vs. "the beginning of the week" vs. "spring"). This again agrees with [15] and [13].
4. Time concepts should serve as time tags for recent episodes. Our *ad hoc* definition of "recent" is roughly the previous week. This has a technical rationale, but the

definition can be changed in future models (in psychology, some would define "recent" differently, e.g. [17, pp. 63]). As episodes get older, their time tags are gradually removed until only the most salient episodes retain their time tags and thus become landmarks [cf. 14, 15].
5. The timing of distant episodes should have a reconstructive nature [14, 15]. Features of episodes that occur regularly should be associated with the time concepts in which they occur (e.g. "toothbrush" – "morning", "snow" – "winter").
6. Time concepts should be both inborn and culturally-based. In computational terms, they should be learnt based on *a priori* knowledge. This seems to be a plausible requirement and it accords with [15].

We propose this list as a set of tentative requirements for timing modules for IVAs. Here, we present a model that works with time periods of weeks, focusing on Points 2, 3, 4, and 6. Concerning hierarchical nesting, we investigated interaction between in-day concepts and in-week concepts. Point 1 has been implemented in our previous model [8] that we further elaborate here, but it has not been reimplemented in the present model. Point 5 is not implemented, although reconstructive timing is theoretically possible in the new model.

III. Architecture and Implementation of Our Agent

We are actually working on several EM subsystems simultaneously and the timing module is only one of them. Each subsystem capitalises on our generic agent architecture (Fig. 1), which resembles universal cognitive architectures by which many IVAs have been inspired. The IVA receives inputs from the *environment* (ENV on Fig. 1) via a simple *attention filter*. The inputs fill up the *perception field* (PF) of the *short-term memory* (STM). The architecture has several memory appendages, such as the *long-term spatial memory* (LTSM) and the *long-term episodic memory* (LTEM). The timing module is part of the LTEM.

A key feature of the architecture is hierarchical decomposition of an agent's behaviour: IVAs' behaviours are decomposed to sub-behaviours, which are further refined until some atomic actions are reached. In fact our architecture distinguishes *tasks* from *goals* making the mechanism resembling the BDI architecture [18], but this is unimportant for present purposes (see [8, 12] for details). Every behaviour may require several *resources*, i.e. objects, for execution. Behaviours to be pursued are selected within the *goal structure* and the *conflict resolution mechanism* based on drives, external events, and a schedule. Currently pursued behaviours are represented within the *task field* (TF) of the STM. The *memory field* (MF) of the STM can hold information about an object recalled from a long-term memory temporarily. Every object is regarded as a tool for action, i.e. it *is* a set of "affordances" [19], meaning it possesses pointers to behaviours it can be used for as a resource (dotted arrows on Fig. 1). IVAs perceive these pointers when observing their environments. The architecture also features an *emotion module*.

Individual parts of the architecture are implemented only when needed, e.g. the IVA with the timing module does not feature emotions. This IVA has been fully implemented using Pogamut 2 [20], a toolkit for fast prototyping of IVAs in Unreal Tournament 2004 [21].

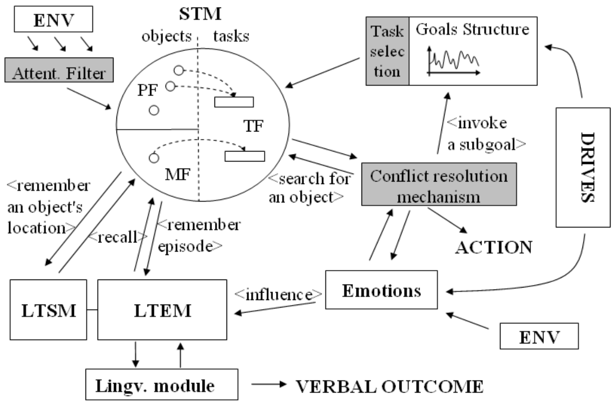

Fig. 1. Our IVA's architecture.

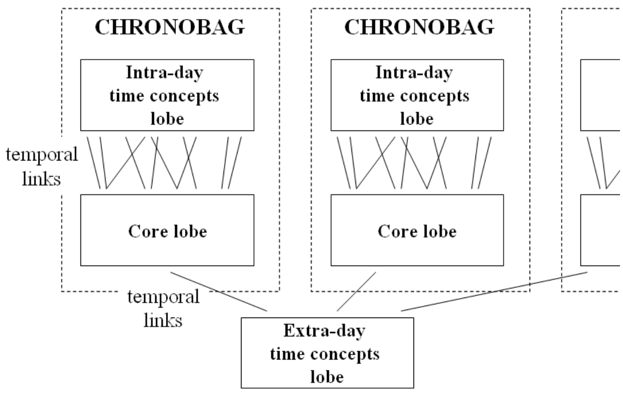

Fig. 2. The structure of one chronobag of LTEM-t.

IV. The Memory Model with the Timing Module

As already said, we extended the previous model we introduced in [8] by the timing module, addressing Points 2, 3, 4, and 6 (Sec. II). The new model is called LTEM-t. It is important to distinguish what has been designed only theoretically from what has been implemented. The features of LTEM-t related to Points 2 – 4, 6 are implemented, but the current implementation of LTEM-t does not have a feature that the previous model had: time pointers (Point 1). Since their reimplementation in LTEM-t would be straightforward and would not bring any new knowledge, we have focused on potentially more promising issues in present work.

Additionally, besides adding the timing module, a conceptual shift has been made while migrating from the old model to LTEM-t. The core module of the old model was conceived symbolically. Currently it is conceived using the connectionist view. This change allows us to address, in the future, the issue of reconstructive recall (Point 5). A reader may perhaps be disappointed that the reconstructive recall has not been addressed yet, but unlike time pointers the reconstructive recall a) as far as we know has never been explored in the field of IVAs, b) is extremely complex and, in our opinion, will demand the research work of several PhD theses to be cracked. Thus, it is far beyond the scope of this paper.

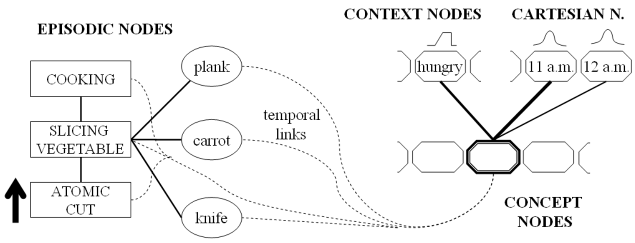

Fig. 3. Left) Hierarchical representation of episodes within the core lobe. For clarity not all episodic links are depicted. Direction of gradual forgetting is denoted by the arrow. Right) The I-lobe's structure. An I-TC node "hungry around noon" and its connections are highlighted.

We now introduce the individual parts of LTEM-t. Note that while the model is psychologically inspired, neuro-/psychological plausibility of its parts is not claimed here. A good starting point for a comparative analysis concerning our model and psychologically-grounded models (which are predominantly not implemented) would be [17, 22].

The LTEM-t structure. LTEM-t is composed of one lobe called the *extra-day time concepts* lobe (E-lobe) and several structures called *chronobags*. Every chronobag consists of two lobes: an *intra-day time concepts lobe* (I-lobe), which is responsible for the representation of time, and a *core lobe*, which is responsible for the actual storage of episodic entries (Fig. 2). The core lobe is the only component present in the old model.

The core lobe. For explanatory purposes assume LTEM-t consists of just one chronobag. The core lobe is a network of interconnected nodes that is built incrementally during storage in the following way: the content of the STM is copied into the core lobe, which entails copying behaviours being pursued by the IVA including sub-behaviours and objects used (but not what other IVAs are doing – that is not represented in the STM; see [23] for the details of our work in progress in that direction). For instance, it could be stored that the IVA is "cooking" (behaviour) and he is currently "slicing vegetables" (sub-behaviour) using a "knife", a "cutting board", and a "carrot" (objects). Every sub-/behaviour and every object is represented as a node in the network. Sub-behaviours are linked with their parent behaviours and with objects used. These links and nodes are called *episodic*. Finally, the network starts to resemble the original hierarchical behavioural decomposition used for the purpose of action selection (Fig. 3/left). The episodic links are weighted, as detailed below. Note that no timing information has been stored yet.

The core lobe stores behaviours of various grain sizes, which allows for *gradual forgetting*: unimportant details of episodes can be "eaten away" from the "bottom" of the network (Fig. 3/left). The IVA can originally remember that he was cooking goulash yesterday morning, including all involved sub-behaviours, but later forget the sub-behaviours, keeping only the high-level information about cooking. Our previous model [8] is likely the only EM model for IVAs having this feature. Additionally that model linked consecutive episodic nodes with time pointers. Recall that LTEM-t does not have the latter feature, but it can be easily extended in this way.

Intra-day time concepts. The I-lobe consists of several kinds of nodes. We will first describe the most important ones: the nodes representing intra-day time concepts ("morning", "after lunch"; I-TC nodes throughout) (Point 2, Sec. II). Think of I-TC nodes as rate-coded neurons with *activity* from $<0, m>$, where m only rarely grows over 3.7, as explained later. Every I-TC node gets activated based on the degree to which the time concept it represents corresponds to the IVA's present situation. Given overlapping time concepts, we have a population coding of subjective time within one day. Assume now that a mechanism for this activation is given.

Assume further that we store information only over one day. Episodic storage *including* timing information works as follows. Let E be a set of episodic nodes being stored in a particular instant and let E be different from the set of nodes that were stored in the previous instant. Every node from E is interconnected with other nodes from E, making a clique of episodic nodes, and *also* with the I-TC node with the highest activity among all I-TC nodes at that moment. The latter link is called a *temporal link* and it represents time information. Both episodic and temporal links are weighted and the "interconnection" actually means the updating of weights:

$$w(t+1) = s(s^{-1}(w(t)) + \gamma c) \qquad (1)$$

That is, if the link already exists between two nodes it is strengthened. Otherwise it is created and its weight is set to 0.5. This value corresponds to the fact that function s is the standard logistic sigmoid $1/(1+e^{-\lambda})$. Its parameter λ is 0.75; this value has been set empirically. The sigmoid assures that weights do not grow beyond 1 and that further strengthening has increasingly lower impact on the weight. γ is the learning rate, set to 1. c is the delta value, set to the constant 0.5. This means that in our present implementation, episodic links and temporal links are updated in the same manner, a simplification we made for the purposes of trialling. In the future c for episodic links can reflect the strengths of entities in the STM and the overall interestingness of the episode (e.g. based on the emotional state of the IVA), while c for temporal links can reflect the activation of I-TC nodes.

Additionally, the winner-take-all (WTA) mechanism for selecting an I-TC node to be linked with the episodic nodes from E can be replaced by a k-WTA mechanism, which would exploit fully the population coding of I-TC nodes. Trialling with that version of the model presents our future work.

The weights are updated *iff* the set E has changed from the last time step or the winning I-TC node has changed from the last time step.

Learning I-TC nodes. Game designers would likely opt for specifying I-TC nodes manually to have the network "under control." However, the I-TC nodes can be also learnt. That is what we did. The advantage is that different I-TC nodes can emerge automatically for IVAs with different "lifestyles" (Point 6, Sec. II).

The I-lobe is actually a 2-layered, feed-forward neural network (Fig. 3/right). The input layer consists of 19 *context nodes* and 24 *Cartesian nodes*; the behaviour of these nodes is our *a priori* information. Cartesian nodes represent objective time. In a biologically more plausible way they can be also conceived as biorhythms, presenting an organism's inborn internal notion of time (on the scale of one day). They have Gaussian-based, weakly-overlapping <0, 1> activation functions. Every hour one of the nodes peaks [see 12 for details]. Context nodes represent external context (e.g. "sunset"), the urgent state of drives (e.g. "hungry"), courses of some activities (e.g. "eating"), ends of some activities (e.g. "after sport"), and emotional states (e.g. "stressed"). Their activation functions are specified by a designer. Biologically speaking, context nodes can be interpreted as representations of internal states that the organism has learnt to recognise, either after birth or evolutionary.

The output layer comprises the I-TC nodes. Weights between the input and the output layer are real numbers from <0, 1> initially set randomly to <0, 0.2>. Equation (2) describes how activation of I-TC nodes is computed; w denotes weights and u the input activation. Equations (3) – (5) are the learning rule, a variant of the Hebbian learning rule with subtractive normalisation [24].

$$v_i(t) = \sum_j w_{i\leftarrow j}(t)\, u_j(t) \tag{2}$$

$$w_{i\leftarrow j}(t\bullet) = w_{i\leftarrow j}(t) + \gamma v_i(t)\, u_j(t) \tag{3}$$

$$y_i(t) = \sum_j (w_{i\leftarrow j}(t\bullet) - w_{i\leftarrow j}(0)) - d \tag{4a}$$

$$x_i(t) = y_i(t) / n \qquad \text{when} \qquad y_i > 0 \tag{4b}$$

$$x_i(t) = 0 \qquad \text{when} \qquad y_i \le 0 \tag{4c}$$

$$w_{i\leftarrow j}(t+1) = w_{i\leftarrow j}(t\bullet) - x_i(t) \qquad \text{when} \qquad w_{i\leftarrow j}(t\bullet) - x_i(t) \ge w_{i\leftarrow j}(0) \tag{5a}$$

$$w_{i\leftarrow j}(t+1) = w_{i\leftarrow j}(0) \qquad \text{when} \qquad w_{i\leftarrow j}(t\bullet) - x_i(t) < w_{i\leftarrow j}(0) \tag{5b}$$

Learning proceeds in two steps. First, the weights are updated according to the basic Hebb's rule (3) [24], where γ is the learning rate. However, each I-TC node can support only a fixed sum of weight increments from the initial weights $w_{i\leftarrow j}(0)$. Thus, second, a subtractive term x_i is subtracted from every weight to prevent weights from outgrowing that sum (4, 5). The total increments are given separately for all weights from Cartesian nodes and all weights from context nodes (d in (4)). The former sum is 2, the latter 1.7, which are empirical values. Thus, while the sum in (2) goes over all input nodes, the sum in (4) goes only over Cartesian nodes and context nodes, respectively, and n is the number of Cartesian (24) and context (19) nodes, respectively. Weights cannot decrease below their initial values (5), which allows for further exploration of the link, helping with re-learning. Due to (5b), it may happen that the total increment is actually larger than 2 or 1.7, respectively, but this discrepancy will tend to be remedied in the next time steps. As a whole the main reason for using this learning mechanism is that it is highly competitive resulting in only about 2-3 context weights plus 2-3 Cartesian weights supporting every I-TC node. The weakly-overlapping activation functions of Cartesian nodes make it likely that consecutive Cartesian nodes will support one I-TC node, which is an advantage. Experiments with various degrees of overlap of activation functions are detailed in [12].

Learnt I-TC nodes capture temporal regularities in the activation of context nodes; the I-TC nodes represent behaviourally relevant time periods. Whereas there will be several nodes for periods during which activations of context nodes change often, like during morning, only one or two I-TC nodes for night may be learnt (because the content of the STM does not change much during night).

I-TC nodes are unnamed; they can be used for internal representation of time but not for communication. To allow IVAs to express themselves conventionally agreed names specified by designers such as "around noon" and "evening" must be assigned to the nodes. A crude natural analogy to this process is a child being taught by its parents to name various parts of day. Note that although the names may be the same for all IVAs, different IVAs will learn different I-TC nodes based on their lifestyles (e.g. early birds vs. late risers). Thus, the internal representation of time may still differ in different IVAs, resulting in the different behaviour of the two memory models (e.g. concerning the accuracy of representation and forgetting rates) and consequently in the different behaviour of the IVAs. An algorithm for the naming is detailed in [12, 13].

Multiple chronobags and the E-lobe. Overall behaviour of LTEM-t depends on the number and behaviour of the chronobags it contains. Chronobags can differ in how many episodes they store and over what periods. Different chronobags can be designed depending on whether one wants to address Point 4 or 5 (Sec. II). Concerning Point 4, a natural tendency is to use one chronobag for representing all episodes that happened during one day, starting every morning with a new empty chronobag. At night every chronobag will be shifted one day into the past, the "today" chronobag becoming "yesterday", etc. During every such shift, i.e. during the night, a selective competition among episodic and temporal links will take place until only the most salient links are kept; these will represent landmarks (this competition extends the gradual forgetting mechanism of the previous model). The E-lobe contains *extra-day time concept nodes* (E-TC nodes) such as "yesterday", "the day before yesterday", etc., each of which is linked with all episodic nodes from the corresponding chronobag.

Another approach, addressing Point 5, is to fill one chronobag with episodes from more days. That chronobag will gradually start to represent an "average day" from the given period, because the night-time selective competition will favour episodes that happen regularly. We can either store just two or three days in one chronobag representing e.g. "the beginning of the week", or a larger period, e.g. "a week". Again nodes of these chronobags will be linked with respective E-TC nodes. It is important to note that an episode may be stored in several chronobags that represent different yet overlapping time periods (Point 3). When different constants for the competitive selection are employed in different chronobags, an episode or its part can be removed from one chronobag but kept in another. Additionally, a kind of "commonplaceness" of events can be determined automatically based on whether an event is well represented in an average-day chronobag. When only uncommon events are stored within single-day chronobags, their timing will become more precise (an earthquake vs. a breakfast). Note that the mechanism of average-day chronobags is reminiscent of the script concept [25].

We have implemented all of these mechanisms.

Recall. For recall the memory structure is not conceived as a neural network but as a spreading activation network. Every question one can ask an IVA, e.g. "when were you swimming?" or "what were you doing yesterday evening?", contains some cues – time, an object, an activity. Nodes corresponding to those cues are injected with activity, which is then propagated to directly neighbouring nodes proportionally to the weight of each link. The most active nodes represent the answer.

V. Implementation and Experiments

The design of the memory system was followed by the successful implementation of a prototype. The prototype was thoroughly tested by a set of experiments [12]. We could test LTEM-t in a maze-style experiment, but we wanted a more complex scenario. Thus, all tests involved a 3D IVA living in an Unreal Tournament 2004 (UT) [21] environment imitating a small town over periods of about a month (one time step represented 15-30 seconds depending on the experiment). By the word "imitating" we mean that the UT map comprised of 6 rooms that were conceived as various parts of the town, i.e. we run the experiments without appropriate graphics. The map featured 27 different objects and places in total.

Summary of the experiments and methods. We tested the memory's ability to learn I-TC nodes, memory robustness, memory space demands, memory accuracy, and episode blending. To test the model under various conditions the IVA lived according to several different lifestyles (student, millionaire, travelling salesman). Lifestyles defined the behaviours to be executed by the IVA and their timing. Every lifestyle usually contained a few regular activities (e.g. breakfast) and a few occasional activities (e.g. going to the cinema). Possible behaviour was represented hierarchically as described in Sec. III. Every IVA had 10 different top-level goals to accomplish (many of them repeatedly) and had the following biological drives: thirst, hunger, weariness, need to urinate and need to wash oneself. The top-level goals were selected to be pursued based on a daily schedule determined by the lifestyle and changes of drives. Note that even the regular activities set out in the schedule started at various times (e.g. the travelling salesman ate breakfast between 7 a.m. and 8 a.m. and the exact time was generated randomly).

We conducted all the experiments using Pogamut 2 [20], which is a toolkit for the fast prototyping of virtual agents inhabiting UT worlds. In every experiment, the I-TC lobe comprised 19 context nodes, 24 or 48 Cartesian nodes, and 40 concept nodes. In general experiments differed in the details of IVAs' lifestyles and in the configuration of chronobags. For brevity, we describe here only the last experiment: see [12] for the others.

Blending episodes. Blending is an interesting phenomenon. It is very natural and common in humans [e.g. 26] and thereby a good candidate for modelling in the field of IVAs. This phenomenon is broad and no-one can claim replicating it computationally in its entirety; however, to our knowledge, reproduction of *any* of its aspects has not been reported yet in the field of IVAs. We detail here conditions under which one aspect can be reproduced in our model.

In this test we asked whether it could happen that the IVA would report that he/she remembered that "he/she was doing X in the evenings last week", but did not remember which days exactly.

Setting. We used the student IVA and simulated him for three weeks. The agent's weekly plan is depicted in Table I. The IVA's LTEM-t consisted of 23 chronobags and respective E-TC nodes: 14 for single days, 7 for pairs of consecutive days and 2 for weeks. Over the first week no episodes were stored but I-TC nodes were learnt (Eq. 2 – 5). Episodic storage (Eq. 1) happened over the last two weeks. In every chronobag the night time selective competition decreased every temporal and episodic weight by 5%. Weights weaker than 0.32 (an empirical value) were removed and their strengths were redistributed proportionally to the remaining weights in the chronobag.

The agent's schedule was generated for every day as follows: 1) biological needs were scheduled (e.g. sleeping, eating, drinking etc.); 2) plans from the weekly schedule were added (e.g. Study, Frisbee, SeePlay for the first day); 3) leisure time activities (entertainment and watching TV) were added to fill the free gaps in the plan. Finally, the agent's daily plan consisted of approximately 10 top-level goals.

TABLE I
AGENT'S WEEKLY SCHEDULE. BESIDES THESE PLANS THE AGENT ALSO HAD REGULAR PLANS SCHEDULED EACH DAY FOR SLEEPING, HYGIENE, EATING, DRINKING, GOING TO TOILET AND FREE TIME ACTIVITIES: ENTERTAINMENT AND WATCHING TV.

	Part of the day		
Day	**Morning**	**Afternoon**	**Evening**
1	Study	Frisbee	SeePlay
2	Study	Frisbee	SeeMovie
3	Study	Frisbee	SeePlay
4	Study	Frisbee, Swimming	SeeMovie
5	Study	Swimming	SeePlay
6	nothing	Swimming	SeeMovie
7	Study	Swimming	SeePlay

TABLE II
EXAMPLES OF AGENT'S GOALS (THE LIST IS NOT COMPLETE FOR BREVITY). SOME GOALS HAVE MORE ALTERNATIVE SUBGOALS FULFILLING THEM (E.G. ENTERTAINMENT), SOME ARE DIRECTLY EXECUTABLE (E.G. SEEPLAY).

Top level goal	Alternatives of subgoals
Entertainment	Read
	DoSport
	PlayComputerGames
Eat	FindResto, EatAtResto
	ShopFood, Cook
SwimmingTraining	GoSwimming, Swim
Study	SitAtTheLecture, GoHome
Frisbee	ConfirmTraining, GetToParkAndPlay
GoToToilet	have no subgoals

TABLE III
ANSWERS FOR A SINGLE-DAY CHRONOBAGS (A) AND AN AVERAGE-DAY CHRONOBAG FOR THE PREVIOUS WEEK (B). IN THE ANSWER COLUMN THERE ARE NODES THAT WERE RECALLED FROM THE MEMORY CUED BY THE "EVENING" CARTESIAN NODES. THE NUMBERS ARE ACTIVATIONS OF THOSE NODES IN THE CORE LOBE. WHEN TWO NODES ARE ON A SINGLE LINE, FOR INSTANCE IN THE CASE OF [EAT - SHOPFOOD], THE FIRST NODE REPRESENTS A TOP-LEVEL GOAL AND THE SECOND A SUBGOAL. IN THIS CASE THE ACTIVATION CORRESPONDS TO THE SUBGOAL. IN THIS PARTICULAR EXPERIMENT, UP TO 5 THE MOST ACTIVE NODES ARE CONSIDERED FOR EACH ANSWER.

Chronobags	Answer
1 day ago (a)	[Eat - Cook, 1.174], [Eat - ShopFood, 1.046 [Eat - Shop, 0.95], [Eat - Carrefour, 0.95]
2 days ago (a)	[Eat - EatAtResto, 1.172], [Eat - FindResto, 1.143], [SeeMovie, 0.832], [DrinkWater, 0.813]
3 days ago (a)	[SeePlay, 1.118], [Drink, 0.98], [Eat - EatAtResto, 0.943], [Eat - FindResto, 0.886]
4 days ago (a)	[SwimmingTraining – GoSwimming, 1.175], [Eat – AtResto, 1.086], [SwimmingTraining – Swim, 0.959]
5 days ago (a)	[Eat - EatAtResto, 1.174], [Drink, 1.174], [Eat - FindResto, 1.089], [DrinkWater - FindTap, 0.934]
6 days ago (a)	[Eat - Cook, 1.292], [Drink, 1.153], [Eat - ShopFood, 1.019], [GoToToilet, 0.966]
7 days ago (a)	[Study - StudyAtSchool, 1.178], [Study - SitAtTheLecture, 0.999], [Eat, 0.803]

Results. After the third week the IVA was asked (a) "What were you doing during the evening on a particular day?" or (b) "evenings last week?". Table III shows an example of an answer. The IVA indeed reported some episodes that had happened regularly only in the (b) answer, e.g. playing computer games (which is randomly scheduled goal that had happened three times this week), while other less regular episodes only in the (a) answer, such as eating in a special restaurant, and some in both cases, e.g. going to cinema. Note that the (b) answer was constructed from the week-chronobag, where different instances of an episode were intermixed. The competition strengthened regularities of these instances and removed details that had happened just once, keeping a common "gist". Similar results occurred when the cue was an event ("When were you playing computer games?").

Note that no subgoals were recalled for the "Last week" answer. Only a more abstract knowledge of the top level goals remained in the memory because top level goals had higher activation than subgoals.

Discussion. The results show that the model is able to partially forget the time information of several similar episodes, leading to one aspect of blending. Of course, blending is more than that [26]. For instance, humans can blend different colours in certain conditions. Bending objects' features is out of scope of the LTEM-t. Humans also make timing-specific errors beyond blending, for instance, they can mistake specific days. While the model can forget whether going to cinema happened on Monday or Tuesday, stating that it happened at the beginning of the week (using average-day chronobag for the beginning of the week), it cannot say Monday instead of Tuesday. Still, to our knowledge, this is the first demonstration that an IVA memory model has the capacity to blend episodes under certain conditions. This paves the way for IVA memory researchers to address more complex blending and false memory phenomena.

VI. DISCUSSION AND CONCLUSIONS

In this paper we have presented a new psychologically-inspired, long-term episodic memory model for virtual agents enhanced by a timing module. The model was prototyped and thoroughly tested. Here we have presented that the model is able to blend episodes of the same class and partially forget timing information. Our other experiments demonstrate that the model mimics also other psychological phenomena, including error proneness of timing at the scale of hours qualitatively similar to those of humans, adaptation to different time zones after a time shift, and usage of different time scales [12].

Is the model ready for practical usage? We will analyze this question from the standpoints of academia and the gaming industry. From the academic perspective several issues remain unaddressed. First, the model works with time periods of weeks, but the real scientific challenge is the entire period of human life. While the I-lobes and the core lobe may remain intact for longer time periods, many kinds of E-TC nodes should be added (and perhaps also learnt). This will bring new questions such as the interaction of multiple timescales, blending over larger periods, and emergence of absolute time concepts ("the week when I got married") and their interaction with relative ones ("last week"). Second, the learning rules the model uses should be compared to other rules. Third, while timing information can deteriorate, the IVA cannot mistake the date. Here a probabilistic interpretation of weights may help. Fourth, the time pointers present in the previous model can be incorporated into LTEM-t. This would represent a first step towards the implementation of reconstructive timing. However, for general reconstructive recall other issues far beyond the scope of timing will need to be addressed, such as the emergence of false memories [26]. A model extended by these features can contribute also to psychology.

From the industry perspective our results suggest that a refined version of the model can be used in many NPCs and IVAs. However, by "refined" we do not mean an extended version but rather a lightened one. From the industry standpoint it does not make much sense to investigate the above-mentioned academic issues unless some of them directly stem from the requirements of a real-world application. Because tuning of the model's parameters is a time consuming process, it makes more sense to simplify the model. For example, I-TC nodes can be learnt off-line or even hard-coded by a game designer. Some technical issues that need to be addressed before the model can be used also exist: most notably consumption of computational resources. Here the model's advantage is that memory requirements can be controlled by forgetting. Also at some point during development a believability study with human subjects will likely have to be conducted.

We presently focus only on some of the academic issues: most notably on false memories and representing actions that happened to other agents including human players [23].

REFERENCES

[1] D. Doherty, C. O'Riordan, "Toward More Humanlike NPCs for First-/Third-Person Shooter Games," in *AI Game Programming Wisdom IV*, Charles River Media, pp. 499—512, 2008.

[2] E. D. Tulving, *Organization of memory*. New York: Academic Press, 1972.

[3] C. Brom, J. Lukavský, "Towards virtual characters with episodic memory II: Episodic memory strikes back," in *Proc. Empathic Agents*, AAMAS workshop, pp. 1--9, 2009.

[4] 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*, Springer-Verlag, pp. 606--617, 2007.

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

[6] W. C. Ho, K. Dautenhahn, C. Nehaniv, "Computational Memory Architectures for Autobiographic Agents Interacting in a Complex Virtual Environment," In *Connection Science* 20(1), pp. 21—65, 2008.

[7] A. Nuxoll, "Enhancing Intelligent Agents with Episodic Memory," PhD thesis, The University of Michigan, 2007.

[8] C. Brom, K. Pešková, J. Lukavský. "What does your actor remember? Towards characters with a full episodic memory," in *Proc. of Intelligent Virtual Agents*. LNCS 4871, Springer-Verlag, pp. 89-101, 2007.

[9] M. Harbers, K. van den Bosch, J. Meyer, "A methodology for developing self-explaining agents for virtual training," in *Proc. of 8th AAMAS*, pp. 1129-1130, 2009.

[10] A. Nuxoll, D. Tecuci, W. C. Ho, N. Wang, "Comparing Forgetting Algorithms for Artificial Episodic Memory Systems," in *Remembering Who We Are - Human Memory for Artificial Agents, AISB,* 2010.

[11] T. Bickmore, D. Schulman, L. Yin, "Engagement vs. Deceit: Virtual Humans with Human Autobiographies," in *Proc. of Intelligent Virtual Agents*. LNCS 5773, Springer-Verlag, pp. 6-19. 2009.

[12] O. Burkert, "Connectionist Model of Episodic Memory for Virtual Humans," Master's thesis, Charles University in Prague, 2009.

[13] O., Burkert, C., Brom, R. Kadlec, J. Lukavský, "Timing in Episodic Memory: Virtual Characters in Action," in *Proc. AISB symposium Remembering Who We Are – Human Memory for Artificial Agents*, pp. 1-8. 2010.

[14] W. J. Friedman, "Memory for the Time of Past Events," in *Psychological Bulletin* 113(1), pp. 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, pp. 129—156, 1995.

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

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

[18] M. E. Bratman, *Intention, plans, and practical reason*. Cambridge, Mass: Harvard University Press, 1987.

[19] J. Gibson, *The Ecological Approach to Visual Perception*. Houghton Muffin, 1979.

[20] R. Kadlec, J. Gemrot, O. Burkert, M. Bída, J. Havlíček, C. Brom, "Pogamut 2 – a platform for fast development of virtual agents' behaviour," in *Proc. CGAMES 07*, 2007.

[21] Epic Games, *Unreal Tournament 2004*. URL: http://www.unreal.com [2.3.2010]

[22] K. A. Norman, G. J. Detre, S. M. Polyn, "Computational models of episodic memory," in *The Cambridge Handbook of Computational Cognitive Modeling*. pp. 189-225, 2008.

[23] R. Kadlec, C. Brom, "I can (almost) remember what you are doing: from actions to tasks," in *Remembering Who We Are - Human Memory for Artificial Agent, AISB,* 2010.

[24] P. Dayan, L. F. Abbott, *Theoretical Neuroscience*. The MIT Press, 2001.

[25] R. Schank, R. Abelson, *Scripts, Plans, Goals and Understanding*. Psychology Press, 1977.

[26] C. J. Brainerd, V. F. Reyna, *The Science of False Memory*. Oxford University Press, 2005.