Sharing Memories: An Experimental Investigation with Multiple Autonomous Autobiographic Agents

Wan Ching HO, Kerstin DAUTENHAHN, Chrystopher L. NEHANIV, René TE BOEKHORST
Adaptive Systems Research Group
Department of Computer Science, University of Hertfordshire
College Lane, Hatfield, Hertfordshire, AL10 9AB, UK
{W.C.Ho, K.Dautenhahn, C.L.Nehaniv, R.teBoekhorst} @herts.ac.uk

Abstract. The overall aim of our work is to develop a generic adaptive control architecture for autonomous intelligent agents. In our previous work we showed how a single agent's survival can benefit from *autobiographic memory*. In the current work we extend this work by introducing a multi-agent context. Also, we investigate autobiographic agents which are able to *share episodic memories* (sequences of perceptions and actions) in the form of 'communications' with others in order to find necessary resources for 'surviving' in an environment. This Artificial Life perspective allows assessment of the potential benefits of such communications in a multi-agent context. A virtual experimental-based approach deals with different conditions on agent communications. Extending our previous work, we introduce an improved control architectures design for autobiographic agents that allows them to communicate episodic memories. Using detailed measurements of the agents' lifespan, this is compared both with groups of purely reactive agents and non-communicative autobiographic agents. Results confirm our previous research hypothesis that autobiographic memory can prove beneficial, indicating increases in the lifespan of an autonomous, autobiographic, minimal agent. The paper presents results from simulation runs varying the following parameters: a) with/without communications of episodic memories, b) varying the number of agents in the environment from one to five, and c) communications under different types of conditions, including presence/absence of *costs* and *motivations*. Results show that in experiments with certain numbers of agents, communicating memories with others can compensate for the negative effect of generally decreasing lifespan due to multiple agent interference. Experiments on specific combinations of motivation and cost conditions also show that, when cost is incurred for communication, agents perform better if they communicate memories only when they are in need of a resource.

1. Introduction

The overall aim of our work is to develop a generic adaptive control architecture for autonomous intelligent agents. More specifically, we are exploring the design space of agent control architectures with *autobiographic memory*. This paper extends the study of autobiographic memory in autonomous agents [1] by studying multiple autobiographic agents in a static virtual environment with different conditions on the communication of episodic memories. For the purpose of this paper, a '*communication*' is defined as *a sequence of events (perceptions and actions) that some agent has acquired in its memory* *through experience and that is being transmitted to another recipient agent*[1]. In this way, an agent can share another agent's episodic memory. Autobiographic memory is a specific kind of episodic memory, which develops in human childhood [2]. Autobiographic agents are agents which are embodied and situated in a particular environment (including other agents), and which dynamically reconstruct their individual history (autobiography) during their lifespan, as defined in [3]. Autobiographic memory is an important ingredient for socially intelligent agents [4]. Moreover, it can be useful for synthesizing agents that can behave adaptively [5], and for designing agents that apparently 'have a life' and thus appear believable and acceptable to humans [6].

From the perspective of ethology and sociobiology, it is increasingly acknowledged that information sharing may improve the performance of animals, e.g. foraging activities in birds [7] and collective behavior shown by social insects [8]. Humans and other primates, such as chimpanzees, live in an even more complex spatial and social field that includes social learning and individualized interactions [9]. As autobiographic agents, humans are constantly telling and re-telling stories about themselves and others to establish social relationships and connectedness: telling (part of) a plausible autobiographical story to others is more than showing a plausible sequence of episodic events, it includes the construction of a plausible story based on one's goals, intentions and motivations [10]. From the perspective of minimal Artificial Life agents we are concerned with a) in what situations communications of memories should occur, b) what kinds of information should be exchanged, c) how costs attached to communications impact on the results, and d) how private or shared episodic information could benefit a single agent and the whole multi-agent system consisting of minimal autobiographic agents.

In our work we are pursuing a bottom-up, Artificial Life perspective, studying communications in the context of an *Event-based* mode of the *Trace-back memory control architecture* derived from our previous work with *single* autonomous agents [1]. We now investigate the effectiveness of this control architecture for the overall survival of *multiple agents* within virtual environments. The environment studied in the present work is based on an ecological environment, i.e. a virtual world that contains a single type of resource ("food") and where agents need to address simple tasks in order to survive (e.g. finding food, avoiding obstacles). Similar environments have been used widely in Artificial Life research.

1.1 Related Work

Studies in psychology have pointed out the cognitive, as well as social, function of autobiographic memory underlying all of human story telling and history-making narrative activities. It has been suggested that the primary function of autobiographic memory is sharing memory with other people [2]. In the context of artificial systems (e.g. robots, agents) research has suggested that the historical grounding of autobiographical agents could in itself be adaptive and help to develop the individualized social relationships and forms of communication which are characteristic of social intelligence [5]. It may also lead to more appealing and engaging 'human-like' interactions with artificial agents, making them more pleasant and acceptable to humans [6].

Previous experimental research on autonomous, autobiographic agents has only

---

[1] All 'communications' in this paper will be receptions of episodic information of this kind. In contrast, ethologists define communication as a signaling by one organism that, when responded to by another, confers some advantage (statistically speaking) to the signaler or its group [11]. Whether and how communications (in our sense) of memories actually do confer some advantage as required by the ethological definition will be a primary concern of our research.

addressed how autobiographic memory might benefit a *single* agent. In a first set of experiments Situation-Action-Situation triplets were used as the core of the agents' autobiography, and the trajectories and lifespan of purely reactive (sensory-driven) and post-reactive (memory-driven) control agents were compared [12]. Results showed that autobiographic memory embedded in the control architecture of the agent can effectively extend an agent's lifespan. A second set of experiments described in [1] studied two separate autobiographic memory architecture mechanisms which are named *Trace-back* and *Locality*. In the Trace-back mechanism, *Event-based* and *Time-based* are two distinct methods for making memory entries, and both of them have been examined in the same static environment with and without noise interference. The autobiographic agents' lifespan in these experiments showed superiority over Purely Reactive agents. In addition, the results also showed that in the static environment without noise interference, an increase in the length of an agent's memory tends to extend the agents' lifespan. Interestingly, when using noise interference, the opposite tendency was observed, i.e. the agents' lifespan eventually decreased as memory length grew too large.

Various memory architectures have been developed in the fields of Cognitive Science, Robotics and Artificial Intelligence, usually for complex tasks such as navigation, case-based problem solving, etc. A detailed review of this work goes beyond the scope of this paper. However, our approach differs as it focuses on bottom-up Artificial Life principles. We study how autobiographic memory and sharing it can benefit an autonomous agent. Consequently, we focus on the effect of sharing memories from an agent-centered perspective and are not concerned with issues such as optimization, path planning, the use of computational techniques such as dead-reckoning, etc. This conceptual background is fleshed out in more detail in [5].

2. Research Hypotheses

Based on the previous experimental framework [1], which studied different possible mechanisms for developing autonomous autobiographic agents from an agent-centered viewpoint, in this paper we investigate the hypotheses that

a) In a multi-agent context, environmental dynamics are complicated by multiple agents wandering around in the same environment. As a result, the interference of encountering other agents is expected to adversely affect the performance of both autobiographic agents and purely reactive agents.
b) Autobiographic agents which communicate memories under appropriate conditions should generally outperform autobiographic agents which do not share memories.
c) If communication of memories has a cost, we hypothesize that communications will be less beneficial than when communication is cost-free.
d) If costs are incurred for the communication of memories, we expect selective strategies for when to communicate to be beneficial to autobiographical agents.

We are investigating communications taking place in different situations, i.e. agents acquire potentially useful information of more efficient paths to the resource either 1) any time they have the opportunity to copy memories of another agent — i.e. they are sufficiently close to another agent, or 2) if, in addition, the agent's internal state is lower than a certain threshold ("hunger"). An important factor which must be taken into account in our experiments is whether communication costs are incurred each time information is transmitted from one agent to another, due to time and energy factors that impact information exchange among physical agents such as animals or robots.

Although simulation environments, in general, can provide an agent with full information about events and features in the environment, we deliberately take a bottom-up, *agent-centered* approach: We do not use any unique object identifiers in the virtual environment so that an agent is not able to identify or recognize *specific* targets. Similarly, the unavailability of global coordinate information of distinct objects restricts the agents' resource tracking system.

3. Experiments

The environmental test-bed is implemented by using the web-based 3D technology VRML, which provides a set of wide-ranging definitions of geometric shapes for building a suitable virtual world with 3D objects. To achieve a certain level of autonomy for virtual agents, time constraint programming techniques [13] are applied to the JavaScript for implementing the control architectures that generate the agents' behaviors. In each time step the script nodes which acts as the 'brains' of the agents route necessary information to modify the shape of both the agents' bodies and the objects in the environment, so that autonomous behavior and continuous movement can be carried out. Variable numbers of autobiographic agents were evaluated for the ability to survive under different experimental conditions relating to communication occurrence (never, only when a potential recipient is hungry, or whenever one agent encounters another one) and whether or not cost is incurred. The environmental resource was not limited so the agents are not in competition. The agents' task was to maintain their energy level within a homeostatic range while exploring the environment and thus not remaining near a resource permanently.

3.1 Virtual Environment & Virtual Agents

In order to assess the impact of communications between autobiographic agents without unnecessary complexity, we designed a simple, large size static environment with irregular object distribution with basic geometric shapes which are generated in real-time from VRML. The virtual environment consists of one resource, four obstacles and walls which serve as the boundaries of the environment. All experimental simulations were running in the same virtual environment in order to allow proper comparisons of the results. The overview of the object distribution in the virtual environment is illustrated in Figure 1.1. Obstacles and walls are static, but the presence of other agents introduces dynamic aspects to the environment.

Basic cylinder and cone shapes are used to construct the body and the sensing area of the virtual agents, as shown in Figure 1.2. The simple design of the virtual environment and the plain appearance of the virtual agent are necessary according to considerations of the performance of current PC hardware, as the speed of the simulations can be noticeably decreased in a multi-agent context. The sensor of the virtual agent is modeled as effectively working within a certain range and angle. Furthermore, we implemented principles of the agent programming language PDL [14], which continually updates sensory and internal quantities using the latest readings from sensors and which then sends new values of the action parameter quantities to the actuators. PDL was originally designed to specifically support dynamical, life-like behavior for autonomous mobile robots.

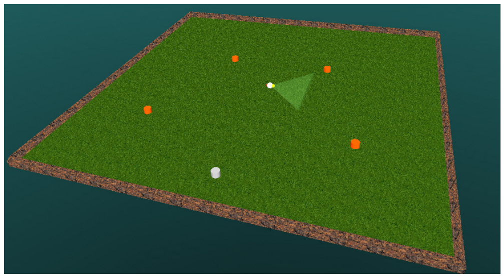

Figure 1.1. Resource and object distribution in the virtual environment. (Cylinder in foreground represents the resource; others represent obstacles.)

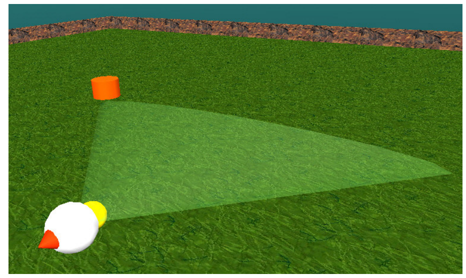

Figure 1.2. A virtual agent and its sensing area (triangle).

3.2 Reactive and Autobiographic Memory Architectures

We mainly study and improve two architectures of agents from the previous work in [1], which are based on a subsumption control architecture [15], namely *Purely Reactive* and *Trace-back*. All agents have a finite lifespan and are required to wander in the environment as their basic behavior. The survival of an agent depends on maintaining homeostasis for the internal state, which is initialized close to a maximum value at the start of each experimental simulation run. Each translation or rotation of the agent will reduce the internal variable by a certain value. When the internal state drops below a threshold, which is half of the maximum value, then the agent begins searching around for the resource located in the environment. If the agent is not able to detect the resource needed, and if the value of a particular internal state is less than a particular minimum value, then the agent will die. The experimental parameters (thresholds etc.) that allow the agents to live in the virtual environment, but eventually die, were determined in initial tests.

Purely Reactive. The architecture of the Purely Reactive agent includes three layers and higher-level behaviors which inhibit or override lower-level behaviors. The agent usually wanders around in the environment by executing the bottom layer in the architecture. At the same time, PDL control for increasing or decreasing the velocity of the agent is employed to achieve a certain level of continuous, natural movement; for example, the agent would slowly increase its velocity after it finished an obstacle avoidance behavior and when it started again to explore the environment. When the agent encounters an object, which can be any kind of resource, obstacle or one of the boundaries of the environment (walls), then the agent avoids the obstacle or the wall by generating a random direction rotating its body. This behavior will also be triggered when the agent encounters a resource object, in case the internal variable associated to the particular resource is higher than the corresponding threshold. Figure 2.1 shows the control architecture for the Purely Reactive agent.

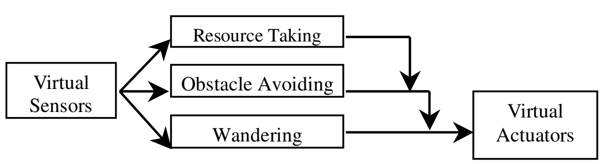

Figure 2.1. Behavior hierarchy which is based on the subsumption architecture for a Purely Reactive agent.

Trace-back. On the basis of the design of a Purely Reactive architecture, Autobiographic agents with memory Trace-back possess an autobiographic memory module on top of the subsumption architecture. An autobiographic agent has a dedicated mechanism for making memory entries as the remembering process, and *using* the memory as a tracing process. In the case of the Trace-back mechanism, the agent has a finite number of memory entries. Introduction of new entries occurs each time the agent experiences an *event*, i.e. encounters either an object or agent, and/or changes its current behavior. This is called *Event-based* memory entry making mode. Each time the agent encounters the resource, all memory entries are cleared and the next memory entry will be made at the first place of the memory table. The index is then set to zero. Each memory entry includes the current direction the agent is facing, the kind of object encountered by the agent (if any), and how far the agent has travelled (distance) since the last event. This information is inserted at the current position of the index into the finite memory table. An energy counter is used for calculating the total cost that would be incurred if the agent were to trace back through the memory entries, undoing remembered actions, while trying to obtain the resource. The value of this counter is increased when a new memory entry is made. A comparison process is then executed immediately between this value and the current internal state of the agent. If the current internal state is lower than the value of the counter, it means the agent would not be able to successfully finish the tracing process and go back to the position of the previously encountered resource, so all entries in the memory table are cleared in this situation. Figure 2.2 illustrates the memory architecture of the Trace-back mechanism. Figure 2.3 shows samples of memory entries made by executing the *Event-based* mode.

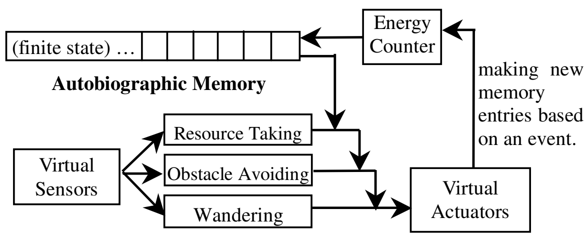

Figure 2.2. Memory architecture of the Event-based Trace-back mechanism.

Index	Direction	Distance
1	1.3525	53.2689
2	-0.5896	0
3	3.0587	15.7635
…	…	…

Figure 2.3 Autobiographic memory sample entries. During the Trace-back process, inverse action sequences are executed for each memory entry in reverse order, e.g. for an agent with autobiographic memory entries as above, to return to the previously encountered resource, Index 1 should be the last action sequence to be undone.

The memory trace back process will be triggered if the internal variable of the agent is lower than the threshold and the table of memory entries is not empty, which indicates that the agent has previously encountered the resource. Once trace-back has started, the agents will simply "undo" all previous behaviors. This mechanism has a close connection to the algebraic notion of *inverse* in mathematics [5]. Thus, the agent will execute the reverse of each action step-by-step starting with the most recent action, using the information specified in Direction and Distance. The trace back process will be completed once the agent has executed actions undoing all memories entries and has reached the target resource. At this moment, the agent will start sensing around for the resource. Note, there are possibilities that the resource is not available at this location since the actual rotation and distance value in each entry might have been slightly distorted by accumulated errors created by the noise during the Trace-back calculation process. As a consequence of these accumulated errors the agent might not be able to finish the Trace-back process, which is terminated if the agent collides with any other object or agent in the environment. Effects of such 'disturbances' are particularly likely in the multi-agent context which could therefore obstruct the *Track-back* process of the autobiographic agents. In a multi-agent context, the same control architectures are running on all agents.

3.3 Agent Communications of Episodic Memories

Effective communication of episodic memories is the fundamental research issue addressed in this paper. Communications can only come from agents which have already acquired an autobiographic memory, i.e. where at least one memory entry has been made.

Communications can be conditioned to occur in two ways: 1) an agent will always receive memory information from a perceived, nearby agent, or 2) an agent perceiving another will receive memories only if it is 'hungry', i.e. if the internal state of the agent is lower than a certain threshold. Nevertheless, based on a general assumption that any memory is more beneficial than no memory, an agent replaces its own memory entries by the information from other agents when either its memory table is empty, or if the action sequences that can be copied from the other agent's memory would take less energy to get back to the resource, thus giving a better path to the resource that can later be used in the Trace-back process. These two conditions suggest that only potentially *useful episodic information* (useful as evaluated from the perspective of the agent copying and later using this information) will be copied.

Communication costs represent a certain value that will be deducted from a recipient agent's internal state for each memory entry being transmitted from one agent to another.

3.4 Design of Experimental Runs

Experiments for examining the survival of autobiographic agents in the virtual environment can be generally classified into three main types according to the use of autobiographical memory and communications: Purely Reactive Agents (*PR*), Autobiographic Agents with no communication (*Mem NoComm*) and Autobiographic Agents with communication (*Mem Comm*). The last type (*Mem Comm*) is again split into two categories that differ in the way the agents communicate: 1) at any time (*Mem Comm*); or 2) only when an agent is hungry (*Mem Comm Hungry*). Each of these two categories can be used in two different conditions: no communication costs attached (*NoCost*), or communication costs attached (*CommCost*). In total, this results in 26 different experimental conditions, each of them tested in 50 runs which take approximately three hours for running on a desktop computer (Pentium 4, 2GHz, 512MB main memory). In order to ensure robustness of the results, in each of the 50 runs the agents were initialized with different internal variables and random locations. The 26 experimental conditions are labeled according to Figure 4.1.

<table>
<tr><th>Conditions</th><th>PR</th><th>Mem NoComm</th><th>Mem Comm NoCost</th><th>Mem Comm Hungry NoCost</th><th>Mem Comm CommCost</th><th>Mem Comm Hungry CommCost</th></tr>
<tr><td>SA</td><td rowspan="5">Fig. 5.1</td><td rowspan="5">Fig. 5.1
Fig. 5.2</td><td></td><td></td><td></td><td></td></tr>
<tr><td>MA_2</td><td rowspan="4">Fig. 5.2
Fig. 5.3
Fig. 5.4</td><td rowspan="4">Fig. 5.3</td><td rowspan="4">Fig. 5.4
Fig. 5.5</td><td rowspan="4">Fig. 5.5</td></tr>
<tr><td>MA_3</td></tr>
<tr><td>MA_4</td></tr>
<tr><td>MA_5</td></tr>
</table>

Figure 4.1. Conditions for experimental runs, whose results appear in figures 5.1 to 5.5. In the first column of the experimental conditions, SA stands for single agent and MA_X for multiple agents where X is the number of agents allocated in the environment for each experimental run.

3.5 Results

Figures 5.1 to 5.5 show five comparisons of different experimental sets. In each figure, the Y-axis represents the lifespan of the agents and the X-axis represents the number of agents in each set of experiments. Each point in the figure is the average lifespan of 50 runs of each agent in a specific experiment; therefore, multiple points appear in the same column of multi-agent experimental conditions. To clearly illustrate these results, each fitted curve in the figures corresponds to the average value of those points in a particular set of result.

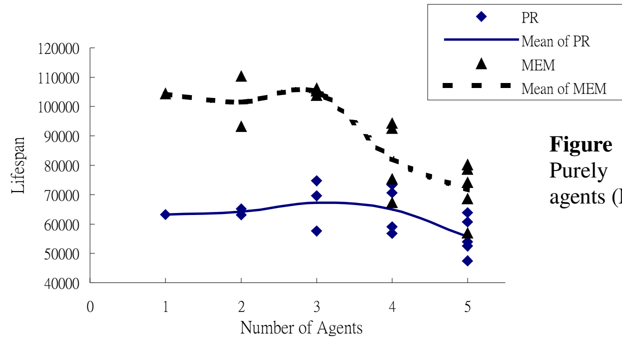

Figure 5.1. Comparing the average lifespan of Purely Reactive agents (PR) and Autobiographic agents (MEM) without communication of memories.

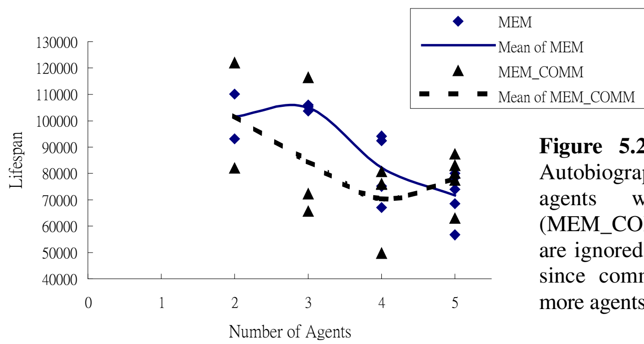

Figure 5.2. Comparing the average lifespan of Autobiographic agents (MEM) vs. Autobiographic agents with communication of memories (MEM_COMM). Results of a single agent in MEM are ignored for the comparison with MEM_COMM, since communication only happens when two or more agents exist in the environment.

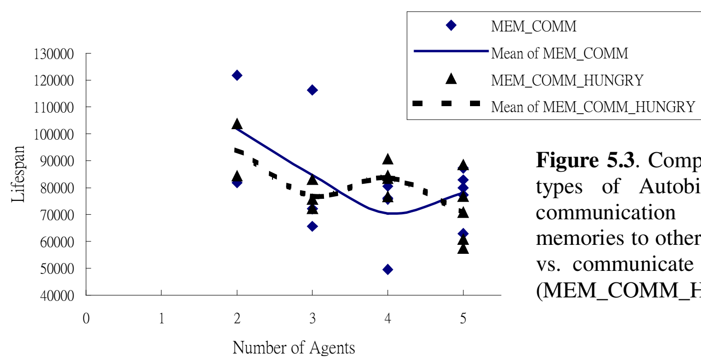

Figure 5.3. Comparing the average lifespan of two types of Autobiographic agents with different communication motivations: communicate memories to other agents any time (MEM_COMM) vs. communicate only when recipient is 'hungry' (MEM_COMM_HUNGRY).

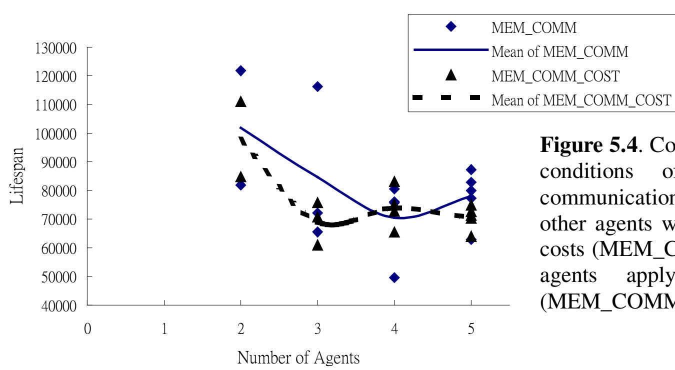

Figure 5.4. Comparing the average lifespan of two conditions of Autobiographic agents with communications: communicate memories with other agents without applying the communication costs (MEM_COMM) vs. communicate with other agents applying the communication costs (MEM_COMM_COST).

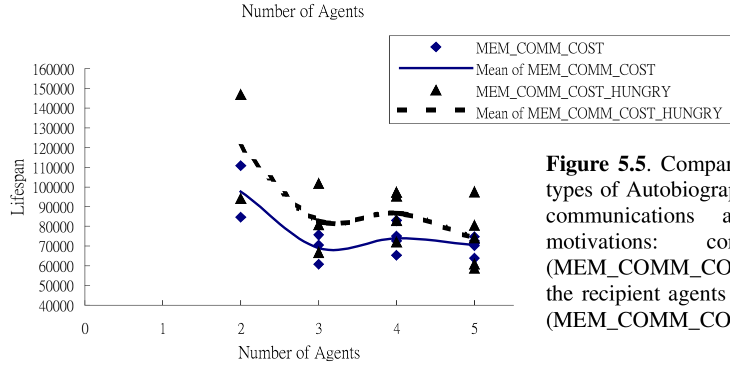

Figure 5.5. Comparing the average lifespan of two types of Autobiographic agents with costs applied to communications and different communication motivations: communication with costs (MEM_COMM_COST) vs. communicate only when the recipient agents are 'hungry' and applying costs (MEM_COMM_COST_HUNGRY).

4. Discussion of Experiments

The experimental results shown in Figure 5.1 reconfirm our previous hypothesis in [1], namely that autobiographic memory effectively extends the lifespan of a Purely Reactive agent which is based on the subsumption control architecture, in both single-agent and multi-agent environments. The agents' average lifespan for both Autobiographic and Purely Reactive agents generally decreases as the number of agents in the environment increases, particularly when the number of agents is more than three in multiple Autobiographic agent experiments, supporting our initial research hypothesis (a). Since more complicated environmental dynamics are created by an increased number of agents, we find higher probabilities of an agent being disturbed by other agents wandering around in the environment during the Trace-back process when there are more agents.

We hypothesized in the multi-agent communication experiments involving the sharing of memories that agents should be able to receive better paths from the current location to the resource. However, no two agents can occupy the same space and have the same perspective on the environment due to their embodiment. This gives rise to the *displacement problem,* the phenomenon that a neighbouring agent's location slightly differs from one's own, which results in a slight error in reusing other agents' paths leading to lower performance, which can be seen in Figure 5.2. This problem reveals the necessity of consideration of embodiment issues in applying memory sharing communication behaviors to Autobiographic agents. One solution to this problem could be to allow agents to correct for the different location and perspective of others. Figure 5.2 also indicates that the increasing frequencies of communications in the multi-agent communication experiments with five agents can compensate fairly well for the negative effect of interference resulting in decreasing lifespans in the experiments of Autobiographic agents without communications.

Figure 5.3 illustrates that there is no significant distinction between the multi-agent communication experiments that directly apply the two different motivational conditions (communication any time versus only when the recipient agents are hungry). Similarly, Figure 5.4 shows no significant difference between Autobiographic agents which incurred costs versus no costs for communicating. Possibly the results would be different if the displacement problem had been solved. Nevertheless, when both factors of communication motivations and costs are combined in the experiments whose results are shown Figure 5.5: there is a tendency showing that when communication incurs costs, agents that communicate with others only when the recipient is hungry perform better than those agents that can communicate any time. This tendency reveals that seeking to use others' episodic memories only when hungry prevents unnecessary costs from being incurred and thus extends the lifespan of Autobiographic agents.

5. Conclusions and Future Work

Our study provides experimental evidence that within our framework autobiographic agents effectively extend their lifespan by embedding an Event-based memory which describes agents' previous action sequences as compared to a Purely Reactive subsumption control architecture. Multi-agent environmental interference dynamics result in decreasing average lifespan of agents. Although we cannot entirely rule out the possibility that various specific conditions of agent communications might be able to improve the overall lifespan of the agents in the environment, some appropriate combinations of factors, e.g. communication motivation and cost factors, result in improved performance. These will be taken into consideration for designing future autobiographic agent architectures and experiments.

With the aim of producing a generic adaptive control architecture with the most efficient usage of autobiographic memory to enhance the survival of autonomous agents, the Trace-back process in the current experimental framework needs to be further improved. For example, Event-based memory making mechanisms inspired by the episodic memory of humans can be manipulated not only by using "undo" in a trace-back process [1], and by re-enacting previous experiences ('day-dreaming' as in [12]), but also by regularly applying consistency checks whereby agents can validate their location, e.g. through recognizing local landmarks such as obstacles, environment boundaries, as well as responding appropriately to unexpectedly encountered object or agents.

Moreover, the experiments reveal the importance of embodiment issues for autobiographic agents; in particular the displacement problem needs to be taken into account when reusing the communicated memories of others. In addition, we are also interested in exploring the use of autobiographical memory in more complex environments, for more complex tasks.

References

[1] Wan Ching Ho, Kerstin Dautenhahn, Chrystopher L. Nehaniv (2003) Comparing different control architectures for autobiographic agents in static virtual environments. *Intelligent Virtual Agents,* Springer Lecture Notes in Artificial Intelligence, Volume 2792, 182-191.

[2] Nelson, K. (1993) The Psychological and Social Origins of Autobiographical Memory. *Psychological Science*, 4:7-14.

[3] Dautenhahn, K. (1996) Embodiment in Animals and Artifacts, *Proc. AAAI Fall Symposium on Embodied Cognition and Action*, AAAI Press, Technical report FS-96-02, 27-32.

[4] Dautenhahn, K. (1999) Embodiment and Interaction in Socially Intelligent Life-Like Agents. In: C. L. Nehaniv (ed): *Computation for Metaphors, Analogy and Agent*, Springer Lecture Notes in Artificial Intelligence, Volume 1562, 102-142.

[5] Nehaniv, C L. & Dautenhahn, K (1998), Embodiment and Memories - Algebras of Time and History for Autobiographic Agents, *Proc. 14th European Meeting on Cybernetics and Systems Research*, Austrian Society for Cybernetic Studies, Volume 2, 651-656.

[6] Dautenhahn, K. (1998) The Art of Designing Socially Intelligent Agents - Science, Fiction, and the Human in the Loop. *Applied Artificial Intelligence*, 12(7-8):573-617.

[7] Templeton, J. L. & Giraldeau, L. (1996) Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behavioural Ecology & Sociobiology*, 38: 105-114.

[8] Bonabeau, E., Dorigo, M., Theraulaz, G. (1999) *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 25-36.

[9] Tomasello, M. & Call, J. (1997) Primate Cognition. *Oxford University Press*, 191-310.

[10] Dautenhahn, K. (1999) The Lemur's Tale - Story-Telling in Primates and Other Socially Intelligent Agents. *Narrative Intelligence, AAAI Fall Symposium.*

[11] Burghardt, G, M. (1977) Ontogeny of Communication. In Thomas A. Sebeok (ed): *How Animals Communicate*, Bloomington, IN: Indiana University Press, 71-79.

[12] Dautenhahn, K. & Coles, S. (2000) Narrative Intelligence from the Bottom Up: A Computational Framework for the Study of Story-Telling in Autonomous Agents. *JASSS, The Journal of Artificial Societies and Social Simulation.*

[13] Saraswat, V, A., Jagadeesan, R., Gupta, V. (1994) Foundations of Timed Concurrent Constraint Programming. *Proc. 1994 IEEE Symposium on Logic in Computer Science.*

[14] Steels, L. (1992) The PDL reference manual. *VUB AI Lab Memo*, 92-5.

[15] Brooks, R. (1986) A Robust Layered Control System for a Mobile Robot, *IEEE Journal of Robotics and Automation*, Vol. RA-2, No. 1, 14-23