How Do Place and Objects Combine? "What-Where" Memory for Human-Like Agents

Cyril Brom[1], Tomáš Korenko[1], and Jiří Lukavský[2]

[1] Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic
[2] Institute of Psychology, Academy of Sciences, Prague, Czech Republic

Abstract. Believable spatial behaviour is important for intelligent virtual agents acting in human-like environments, such as buildings or cities. Existing models of spatial cognition and memory for these agents are predominantly aimed at issues of navigation and learning of topology of the environment. The issue of representing information about possible objects' locations in a familiar environment, information that can evolve over long periods, has not been sufficiently studied. Here, we present a novel representation for "what-where" information: memory for locations of objects. We investigate how this representation is formed and how it evolves using a simplified model of a virtual character. The behaviour of the model is also compared with behaviour of real humans conducting an analogical task.

1 Introduction

Humans act in space. They furnish the space they live in with objects. Be it a van Gogh's painting or a pen, it pays to remember where ones' belongings are located. Intelligent virtual agents (IVAs) usually act in space as well. This demands them to have similar "what-where" information as humans do. In many today applications, IVAs read this information directly from the world map, which corresponds to complete environmental knowledge. While this approach may suffice for static words or dynamic but fully observable words, it results in unbelievable behaviour in dynamic words that are not fully observable. The latter kind of environments is increasingly employed by today applications. For instance, think of non-player characters (NPC) from a role-playing game (RPG), or virtual/robotic companions required to orient themselves in humans' houses. Here, a better approach is needed.

Why not to memorise all objects that an agent encountered? Assume we have a large, partially observable environment with objects that are passive but whose locations can be changed by external forces beyond the agent's capabilities. For instance, a pen can be moved by a fellow agent unbeknown to our agent. In this situation, we can expect multiple memory records of a position of each object based on the history of the object's moves. Where is the pen: at the working table or next to the TV? A simple list or a stack of memory records can not answer this.

How can we improve the performance? Consider objects that humans use. Humans have some organisation of the placement of their belongings; things are not placed randomly within their surroundings, instead, they are clustered purposefully according

Zs. Ruttkay et al. (Eds.): IVA 2009, LNAI 5773, pp. 42–48, 2009.

to their needs and cultural norms. Some objects appear regularly at some places (newspapers inside a mailbox). Other objects are almost never being relocated (a van Gogh's painting). Yet others are being relocated so often that it is not practical to remember their exact position. Consequently, when a human searches for an object, often, a sort of stimulus-response mechanism is employed. For a different object, several places are inspected in a specific order; sometimes, the whole house is scrutinised but starting at a specific place.

This brings us to the notion of *searching rules*, which are basically a sequence of places that should be inspected when searching for an object of a particular kind. Importantly, we mean by place any logically coherent space abstraction. These abstractions can differ in size and can be hierarchically nested; e.g. a bedside table, a place between this table and the bed, a corner of the living room, a living room, a house, etc. This corresponds to the way humans are supposed to cluster space [e.g. 7]. A "what-where" memory for IVAs acting in large, partially observable environments with movable objects should have the ability to develop searching rules.

In this paper, we present a model with this ability. To its advantage, it is quite simple. The model has been integrated with our agent with general memory capabilities [2] and investigated in scenarios mimicking a situation of a person moving into a new house. Specifically, we focused on the questions of how the initial representation is formed, how (and whether) searching rules emerge, and how the model relearns, measuring effectiveness of the model' behaviour. We also investigated behaviour of the model with different parameters, aiming to find a point of "optimality." Additionally, we conducted a simple experiment with human subjects that mimicked the agent's task using a simple Flash application, and evaluated the model qualitatively against the acquired human data.

The results allow us to conclude that the searching rules emerge easily and quickly and that they are qualitatively similar to searching rules developed by humans in their version of the task. The model also relearns well. Additionally, humans' data helped us to isolate one feature of human behaviour that cannot be explained by the notion of searching rules straightforwardly, but it can be added onto a top of the model easily. In overall, our opinion is that the model is now ready to use in real-world applications requiring plausible "what-where" memory.

The structure of the paper is as follows. Sec. 2 reviews related work. Sec. 3 details the memory model. Sec. 4 reviews and discusses the experiments.

2 Related work

From the psychological perspective, the ability to locate objects is a faculty of spatial cognition, which is tightly connected to spatial memory. Spatial memory is conceived as a set of multiple interconnected systems rather than a monolithic block [4].

Recently, several general memory frameworks for IVAs were presented [2, 8, 9, 13, 14]; however, the degree to which they address any issue of spatial cognition is minimal. Fortunately, several works directly focussing on some aspects of spatial cognition for IVAs have emerged to increase believability of their spatial behaviour, including mapping, localisation, and navigation. Noser et al. developed IVAs learning topological structure of the environment and navigating using this structure [12]. In a psychologically more plausible manner, Thomas and Donikian [18] addressed similar issue. A mechanism for anticipating position of an object that can move itself, e.g. a sheep, was presented by Isla and Blumberg [10]. Unfortunately, neither of these works addressed sufficiently the issue of "what-where" memory for passive, but movable objects. Unlike these models, the mechanism of Strassner and Langer [15] directly aimed at representing both topological as well as "what-where" information and this information could gradually deteriorate when not refreshed. However, neither this model was designed to cope with objects that can be moved several times, lacking the ability to develop searching rules.

The field of gaming AI predominantly addresses believable and efficient path-finding and automatic construction of space representation [e.g. 6, 16]. To our knowledge, the issue of "what-where" memory is not addressed. Finally, we are not aware of any work either from computational psychology or robotics that could suit our purpose directly. Robotics tend to focus on the issues of localisation and terrain mapping [e.g. 11], which are "low-level" from the perspective of IVAs, pointing to the significant difference between application domains. Even designers of robots with episodic memory systems [e.g. 5] do not seem to consider the "what-where" issue as a crucial one for their discipline. Psychological experiments investigating spatial abilities of humans are of considerable interest to the field of IVAs [see 4, 17 for reviews of some], but computational psychology tends to produce special-purpose models replicating data gained during laboratory tasks [e.g. 1]. It is hard to imagine a meaningful application in which an IVA could be engaged in such a task directly. Additionally, regarding "what-where" information, psychology tends to investigate what would be called in our context short-term representations of positions of static objects and their mental rotations, e.g. for the purposes of elucidating the allocentric—egocentric tension [reviewed in 4].

3 Model

The model we have developed is a simple associative network. It is composed of two kinds of nodes: *objects* and *places* (Fig. 1). Place nodes represent places with different levels of complexity and they are hierarchically nested. Object nodes have weighted links with place nodes; these stand for "what-where" information: a possible occurrence of a particular object at a particular place. Now, if an object is found by the agent, or comes to the agent's attention, the links to *all* the locations where it has been found are strengthened (e.g. the links from the glasses to a) the bedside table, b) to the bedroom, and c) to the whole house – see Fig. 1).

How could searching rules possibly emerge from this network? Note several things. (1) Links from an object node to nodes representing places at a similar level of complexity approximates the probability distribution of finding the object at given places. (2) Links to nodes representing more abstract places are strengthened more often than links to nodes of concrete places ("glasses are always in the house, but only sometimes at the bedside table"). Now, if queried for an object position, we can find the appropriate object node and past locations of the object's occurrence via the object—place links. Assuming that the pattern of the object's movement will be same in the future as it was in the past, we can arrange the place nodes in order of the strengths of the links leading to them (Point (1)). However, because of Point (2), before we do this, we need to scale the links' strengths by an inverse function of the complexity of places; otherwise, the abstract places will be always first in the list.

The fundamental assumption is that with an appropriate scaling function, the result will be a searching rule, i.e., a list with balanced ordering of concrete and abstract places where the object can be looked for. Concrete places should be first on the list provided that there are only a few concrete places where the object can be found. Otherwise, an abstract place should be first or very close to the beginning of the list. Particular places can be searched directly (e.g. the bedside table) while abstract places should be inspected (e.g. scrutinise the kitchen). The fact that this mechanism really produces believable searching rules is demonstrated in the next section.

An important question is how to deal with distances. The model ignores distances, conceiving the searching rules as verbal answers on the question: "Where do you think is an object X?" Our opinion is that in a middle-sized environment, e.g. in the ground-floor house used in our experiments, we can ignore distances letting the agent search for at most probable places, but distances become important in larger worlds, such as in multi-floor buildings or cities. How to take distances into account? One can take outputs of our model as inputs for an engine solving the *travelling salesman problem* (TSP) with *uncertainties*. Ideally, the engine should reflect how people solve the same problem. Surprisingly enough, it seems the TSP with uncertainties, as opposed to the classical TSP, has not been investigated in psychology until recently [19], and this work is not conclusive from the standpoint of IVAs. Thus, we assign the question of plausible penalisation of distant places as future work.

Formal definition. Formally, the network is a triple $<P, O, E>$.

P is the set of all *place nodes*, each of which is a quintuple $<p$, *up*, *down*, *level, size*$>$, where p is the node, *up* its super-location, *down* its sub-locations, *size* is a number of its sub-locations, and *level* is the level of abstraction. Abstractions are numbered from the bottom: *level* for a specific place is 1, and then the levels are enumerated by one towards the root of the hierarchy.

O is the set of all *object nodes*, i.e. the object records.

E is a set of weighted *edges*, each of which is $< x$, *found*, *missed* $>$, where $x \in P\times O$ is the edge, and *found* $\in \mathbf{N}$ is the number of times the object was found (and taken) or seen at the particular place, and *missed* $\in \mathbf{N}$ is the number of times the object was being searched for at the place but not found there.

In the present version of the model, we assume that P is specified in advance by a designer and fixed during the simulation (but see also Sec. 4).

Learning. When the model stores positional information about an object, it first checks whether the corresponding object node exists and creates a new one if needed. Then, the *found* variable of links between the object node and place nodes of *all* the places where the object was found are increased by 1 (e.g. the house, the bedroom, the bedside table). When the agent is looking for an object at a specific place and this object is not there, the *missed* variable is increased by 1. If the agent is searching in a location and the object is not in any of its sublocations, *missed* variables for this location and all the sublocations are increased by 1.

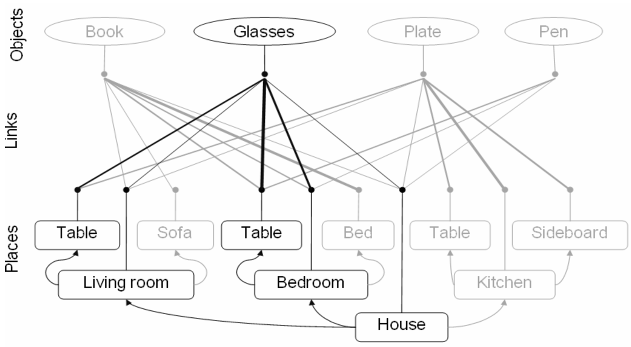

Fig. 1. Spatial memory. Some nodes from the experimental scenario are schematically depicted; the glasses' links are highlighted. The width of the links denotes the size of the *found* variable.

Rules formation. When the model is queried for an object position, *size-normalised trustfulness* (SNT) is computed for every place the object has an edge with as:

$$SNT = f\,(found, missed)\;/\;complexity\;(level, size) \tag{1}$$

The function f determines influence of the *found* and *missed* variables on the estimation of likelihood of finding the object at the given place and *complexity* is the scaling function. There are more options how to choose f and *complexity*. In our experiments, we have investigated how the model behaves with the following functions, where a, b, and c are parameters having been found during the trialling:

$$f = b.found - missed \tag{2}$$
$$complexity = level^{a} \tag{3}$$
$$complexity = size^{c} \tag{4}$$

The places are arranged in order of SNT. More specific places after their super-location are dropped from the rule. The sorted list presents a searching rule.

Notice that (3) completely disregards sizes of places, leading to assigning the same SNT to two rooms: one very small and another very large. Clearly, a human would prefer to inspect the latter room first. This motivates Eq. (4). Further discussion on how f can look like can be found in [3]. Note also that thanks to the equations above, the learning works in a Hebbian manner.

4 Implementation and Experiments

The model has been integrated with our generic agent with episodic memory abilities [2], becoming the agent's the long-term memory for positions of objects (LTSM). The agent also possesses a BDI-based *action selection*, an *emotion module*, an *attention filter* through which only some percepts can pass, a simple *short term memory* – an intermediate stage for object records that are to be later stored in the LTSM –, and an *autobiographic memory* with plausible timing of events and forgetting.

We investigated on a rigor basis the following questions: Do searching rules emerge? Does the time to emergence depend on the frequency of usage of an object?

Do the rules emerge quickly and is searching effective? How quickly the network relearns? Is there one optimal setting for all situations? Are the rules similar to those used by humans; are they believable? The methods and results are detailed in [3]. Here, we present a summary of the methods and discuss the main findings.

For testing the model, we have used a "new house" scenario, where we simulate the agent living for several weeks in a house with 6 rooms to which it just has moved. We defined five objects' classes modelling prototypical behaviour of five distinct kinds of objects, such as "90% class" (there is 90% chance that the object is located at a particular place, and 10% chance that it is located elsewhere; this applies e.g. for a can opener), or "3x30%" (there are 3 places in a same room in each of which the object can be located; 10% chance that it is elsewhere; e.g. glasses). Although we have a 3D implementation of the agent, for the experiments here, we use a 2D world. For validating our model against human data, we developed a simple Flash-based application in which we investigated behaviour of human subjects in a task similar to the task used in experiments with the model (however, note a limitation here: this application tested subjects' memory for time intervals of minutes, not weeks).

In a sum, the searching rules indeed emerged and the time to emergence depended on frequency of object's usage. Most searching rules emerged in less than 6 searches and the subsequent searching was nearly optimal. Additionally, in most cases, searching rules of the model were more effective than humans' behaviour. In general, the model driven by Eqs. (1), (2), and (4) performed better than the model driven by Eqs. (1), (2), and (3). The latter variant is not robust to changes of sizes of rooms.

Two problematic issues were revealed. First, we conducted many variants of the experiment and it was not possible to find one common parameter setting that would suit well for all the variants. The best common values found, that is $a = 9.6$, $b = <1, 15>$, produced behaviour that was worse (summing over all experiments) than average human behaviour by 56% and worse than the best distinct parameter settings by 89%. Especially, for the common parameter setting, there were problems with relearning: the model took it long to forget over-learned places (e.g. after changing a PC to another room). Yet the model relearned well for distinct parameter settings, pointing to the necessity of extending the model with a mechanism deriving automatically parameters based on the context, and perhaps to use another kind of Hebbian rule to alleviate forgetting abilities of the model.

Second, qualitative comparison of searching rules of the model with those of humans revealed that humans behaved differently. However, we identified a simple cause: humans consistently used an additional heuristic to search at a place where the object was found last time. Only after the object was not found there, the subjects turned to their searching rules. When we added this heuristic to the model, the model's searching rules were qualitatively similar to those of humans.

5 Summary

The experiments we conducted showed that, in a middle-sized environment, searching rules emerge easily and the searching for objects is effective and comparable to the searching conducted by a human. However, there is a room for improvement concerning relearning. Our present work is organised around these points: a) testing the model in a larger environment, in particular in a city, b) improving forgetting, c) development of a mechanism that would learn space abstractions automatically. Should the model be used in a real application with a larger world, an engine solving travelling salesman problem with uncertainties should be added. The extended version of the paper detailing the experiments and supplementary videos are available [3].

Acknowledgement. This work was partially supported by the Program "Information Society" under project 1ET100300517, and by the Ministry of Education of the Czech Republic (Res. Project MSM0021620838).

References

1. ACT-R research group. Publications on Spatial Reasoning and Navigation (April 19, 2009), http://act-r.psy.cmu.edu/publications/index.php?subtopic=23
2. Brom, C., Pešková, K., Lukavský, J.: What does your actor remember? Towards characters with a full episodic memory. In: Cavazza, M., Donikian, S. (eds.) ICVS-VirtStory 2007. LNCS, vol. 4871, pp. 89–101. Springer, Heidelberg (2007)
3. Brom, C., Lukavský, J.: Tech. Rep. January 2009. Dept. of Software and Computer Sciences. Charles University in Prague (2009), http://artemis.ms.mff.cuni.cz//main/tiki-publications.php
4. Burgess, N.: Spatial memory: how egocentric and allocentric combine. Trends Cognitive Sciences 10(12), 551–557 (2006)
5. Dodd, W.: The design of procedural, semantic, and episodic memory systems for a cognitive robot. Master thesis. Vanderbilt University, Nashville, Tennessee (2005)
6. Hamm, D.: Navigation Mesh Generation: An Empirical Approach. In: AI Game Programming Wisdom, vol. 4, pp. 113–124. Charles River Media (2008)
7. Hirtle, S.C., Jonides, J.: Evidence of hierarchies in cognitive maps. Memory & Cognition 13(3), 208–217 (1985)
8. Ho, W., Dautenhahn, K., Nehaniv, C.: Computational Memory Architectures for Autobiographic Agents Interacting in a Complex Virtual Environment. Connection Science 20(1) (2008)
9. Ho, W., Dias, J., et al.: Agents that remember can tell stories: Integrating Autobiographic Memory into Emotional Agents. In: Proc. of AAMAS, ACM Press, New York (2007)
10. Isla, D., Blumberg, B.: Object Persistence for Synthetic Creatures. In: Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (2002)
11. Kuipers, B.: The Spatial Semantic Hierarchy. Artificial Intelligence 119, 191–233 (2000)
12. Noser, H., Renault, O., et al.: Navigation for Digital Actors based on Synthetic Vision, Memory and Learning. Computer Graphics 19(1), 7–19 (1995)
13. Nuxoll, A.: Enhancing Intelligent Agents with Episodic Memory. PhD thesis, The University of Michigan (2007)
14. Peters, C., O'Sullivan, C.: Synthetic Vision and Memory for Autonomous Virtual Humans. Computer Graphics Forum 21(4), 743–752 (2002)
15. Strassner, J., Langer, M.: Virtual humans with personalized perception and dynamic levels of knowledge. Comp. Anim. Virtual Worlds (16), 331–342 (2005)
16. Sturtevant, N.R., Buro, M.: Partial Pathfinding Using Map Abstraction and Refinement. In: Proceedings of AAAI, pp. 1392–1397 (2005)
17. Taylor, H. A., Rapp, D. N.: Updating Human Spatial Memory. In: Animal Spatial Cognition (2006) http://www.pigeon.psy.tufts.edu/asc/taylor/ (April 19, 2009)
18. Thomas, R., Donikian, S.: A spatial cognitive map and a human-like memory model dedicated to pedestrian navigation in virtual urban environments. In: Barkowsky, T., Knauff, M., Ligozat, G., Montello, D.R. (eds.) Spatial Cognition 2007. LNCS (LNAI), vol. 4387, pp. 421–438. Springer, Heidelberg (2007)
19. Wiener, J.M., Lafon, M., Berthoz, A.: Path planning under spatial uncertainty. Memory and cognition 36(3), 495–504 (2008)