Interacting with a Robot to Enhance Its Perceptual Attention

Yuval Marom and Gillian Hayes
Institute of Perception, Action and Behaviour
Division of Informatics
The University of Edinburgh
5 Forrest Hill, Edinburgh EH1 2QL, UK
{yuvalm, gmh}@dai.ed.ac.uk

5.4.2001

Abstract

Robots are like humans in that they cannot and do not pay attention all the time; they do not process all the information perceived, and make decisions about what is worth their attention. If we put a robot in a social environment (either with other robots, or humans), this can help it analyse perceptual information: modelling attention is very hard, and a social context simplifies the problem. The purpose of this paper is to examine the effect of interacting with a robot on the characteristics of the data it is exposed to. We want to show that the perceptual data is better structured when the robot is in a social situation. We analyse the data using Principal Component Analysis, which is a multivariate data-analysis tool. We do not actually present here a computational attention algorithm for analysing perceptual data, although we will mention the system we are currently developing, which is based on the Self Organising Feature Map.

1 Introduction

We do not pay attention all the time, and do not process all the information we perceive. Robots are the same. We make decisions about what is worth our attention based on internal factors, and on interactions with both the physical and social environments.

Most of the work on attention in Computer Science and Artificial Intelligence has mainly focused on the spatial dimension of the problem: given a snap-shot of the environment, find out *where* are the locations worth attention for further processing. However, for an on-line system such as a robot, we believe the temporal problem is equally important: *when* should a snap-shot be taken?

1.1 Attention

By attention we are referring to the interaction of structures in memory in response to stimuli, rather than one specific mechanism. In other words, we are interested in the way that perceptions influence the dynamics that occur in memory. For example, how structures in memory habituate to familiar stimuli, and exhibit an orientation response to novel ones (Cowan, 1988; Posner et al., 1980; Sokolov, 1963; Näätänen, 1992).

Our motivation stems from the need to equip a robot that exists in a real, physical environment with an attention system that 'analyses' perceptual input. Robots need to deal with perceptual overload, and consider memory limitations. However, the real world is perceptually very complex and noisy, and we believe it is worth trying to simplify the problem in order to be able to focus on the mechanisms involved. We believe a significant simplification is achieved when the robot is placed in a social context.

Proc. TIMR 01 - Towards Intelligent Mobile Robots, Manchester 2001.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-01-4-1. http://www.cs.man.ac.uk/csonly/cstechrep/titles01.html

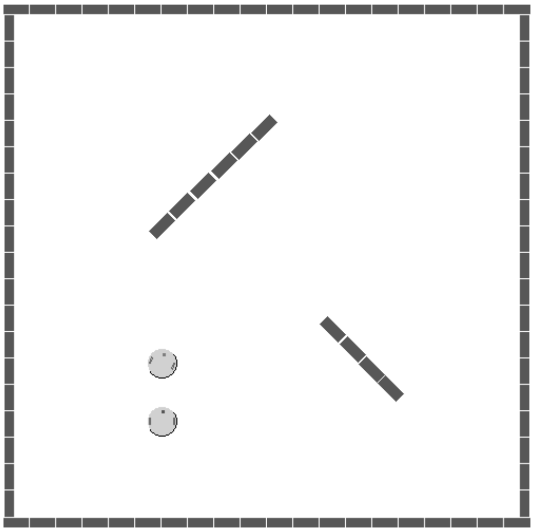

Figure 1: The simulated environment.

1.2 Social Interactions

Social learning has been divided into different sub-categories, depending on what level of information is assumed to be transmitted (for an overview on Social Learning, see Zentall and Galef Jr. (1988); Whiten (2000)). A good analysis of different levels of information transfer is given by Conte (2000), who distinguishes between *social facilitation* and *imitation*. In social facilitation a model merely acts in ways to increase the chances that the observer perceives given events, whereas in imitation information about goals is transmitted as well. We will deal with the former type of social learning, which is similar to what has been termed *stimulus enhancement* in the animal literature (Spence, 1937).

We refer to a 'teacher' interacting with a 'learner', to enhance its ability to analyse perceptual information. There are different kinds of interactions possible in this relationship, ranging from the learner passively following the teacher in an environment, through 'bodily' interactions such as the teacher pointing to objects (Scassellati, 1998) or positioning the learner, to full '2-way' caretaker-infant interactions where the teacher waits for feedback from the learner and vice-versa (Breazeal and Velásquez, 1998). In this paper we will deal mainly with the most simple form — 'following': the teacher leads the learner through the environment, and the learner follows behind.

From a more Engineering point of view, if a robot could learn by interacting with a human (or with an experienced robot that is different morphologically and architecturally), we might avoid explicit programming and parameter-tuning.

2 Perceptual Analysis

In this paper we will *not* actually present an attention system, although we will mention it in Section 4. Rather, we will examine here the value of interacting with a robot in enhancing its *ability* to analyse perceptual data. We will record and analyse datasets acquired by sensory inputs under different conditions, with and without a teacher. The method we will use for the analysis is *not* suggested as the actual attention algorithm for the robot, although it could well serve as such. See more on this in Section 4.

Our aim is to show that the quality of the robot's experience is improved in a social situation, in the sense that the information available to it for further processing (for example learning) is more structured than when it wonders around on its own.

In the following experiments we are looking at one task, wall-following, in both a simulated and physical environment. The simulated experiments are performed using a Khepera simulator with a learner agent following behind a teacher agent, both using infra-red sensors (see Figure 1). The 'data' here come from 6 sensors around the front of the learner. The physical experiments are performed using

Figure 2: The B21 robot following a human. The actual size of the arena is approximately a 3m by 4.5m square

a RWI-B21 robot and a human; the robot can detect and follow a human using its on-board video camera (see Figure 2), and it can also be controlled externally using a joystick. The 'data' here come from 20 sonar sensors around the robot. The actual size of the arena is approximately a 3m by 4.5m square.

The exploratory investigations we present here are carried out under these conditions: (a) the robot is experienced about the task (hand-crafted behaviour), (b) an experienced demonstrator interacts with the robot (either through a following behaviour, or with the joystick), and (c) the robot is inexperienced and on its own (random behaviour). In the simulation, 'following' is achieved by the teacher performing the hand-crafted behaviour, and the the learner following behind it using a built-in following behaviour. The B21 robot follows a human demonstrator around the arena using its video camera (the visual input is *only* used for tracking purposes); in a separate experiment, the joystick is used to drive the robot around the arena. Under these different conditions, we record sensor readings at every step, and analyse the dataset comprising all the sensor readings up to that point, using Principal Component Analysis (PCA) in Matlab.

PCA is a multivariate data-analysis tool used to find where most of the variance in the data occurs, when the dimensionality of the data is high, and one is interested in finding a lower dimensional representation of the data. This is done by finding a linear transformation of the data such that the variance along the new dimensions is maximised. The dimensions of this lower representation are referred to as the principal components of the data, and one can look at the relative variance that the first k principal components contribute to the total variance, as an indication of how well these k dimensions describe the data. In other words, we can examine how well the data can be reduced to k dimensions, where k is some low number.

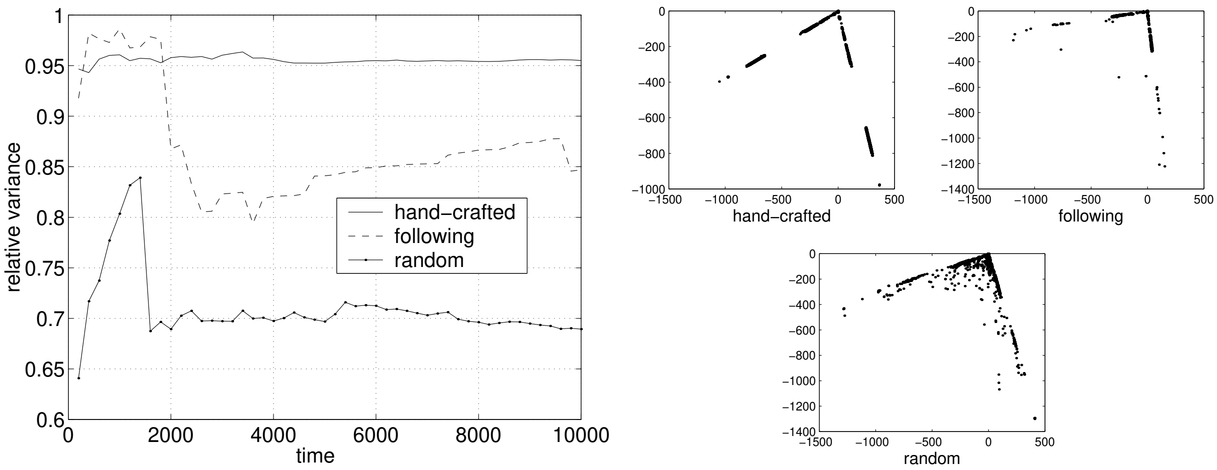

Figure 3: Simulation IR data. Left: convergence of the relative variance of the first **2** principal components. Right: projection of the data onto 2 dimensions, as determined by the first 2 principal components; for visual purposes only a (random) subset of the data is plotted.

In our analysis, we look at the ability of PCA to reduce the dimensionality of the data, over the duration of a run, where a run in the simulation consists of 10000 steps, and with the B21 robot approximately 5000 steps. We expect that as the robot is seeing more data, this ability will change (and hopefully improve). At equal intervals of the run we inspect the proportion of the variance accounted for by the first k principal components, out of a total of n, where n is the dimensionality of the data determined by the number of sensors. At the end of each run, we also look at the representation of the data in a low dimension, after all the data have been accounted for. We do so by projecting and plotting the data in 2 dimensions, as specified by the first 2 principal components. For visual purposes we only plot a small, random, subset of the data.

For a simple task such as wall-following, we would not expect that more than 2–3 principal components would be required. Figure 3 shows the relative variance of the first **2** principal components (PC's) from the simulated experiments (*i.e.* $k = 2$), and the projection of the data onto these 2 PC's. We see that the hand-crafted behaviour exposes the robot to perceptual information where 2 principal components explain 95% of the total variance, compared to 70% by the random behaviour. When the robot follows behind an experienced teacher, the data it is exposed to are explained better than the random behaviour and worse than the hand-crafted behaviour, as expected; further, convergence occurs later than the hand-crafted behaviour, but quite early in the episode (approximately after a fifth of the time).

In the low dimensional representation of the data (right side of Figure 3) we can see quite clearly three clusters for the 'hand-crafted' and 'following' plots. One cluster is the intersection of the 2 apparent 'lines', which is the area of low detection, *i.e.* 'no wall', and as we move away from this area, we reach one cluster on the bottom right corner, and one on the top left, corresponding to the 'right wall' and 'left wall', or vice versa. This means the data are perceptually separated for the cases when the robot is seeing a wall on the left, on the right, and none. In contrast it is harder to identify the clusters in the random plot.

Figure 4 shows the relative variance of the first **3** PC's from the physical experiments (*i.e.* $k = 3$), and the projection of the data onto a lower number of dimensions (for visual purposes the data are projected onto the first 2 PC's, rather than the first 3). Overall, the data in the physical experiments are not as 'clean' as in the simulations, as we expect since there is a lot more noise and many more sensors are used. In the best case, 3 PC's can only explain 60% of the variance (we could get this figure higher by using more PC's, for example the first 10 PC's account for 85% of the variance when the joystick is used). Comparing the results between the different conditions, our conclusions are the same as in the simulations, except here we have an additional scenario: where the robot is guided with a joystick — a much more stringent form of interaction; in this scenario, the robot's perceptual information is enhanced

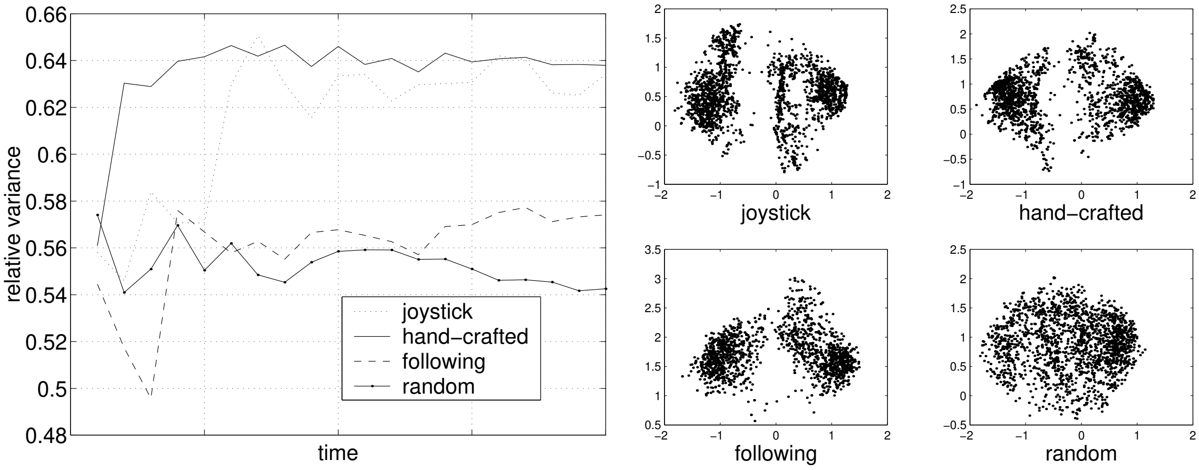

Figure 4: B21 sonar data. Left: convergence of the relative variance of the first **3** principal components; the reason the time labels are omitted from the plot is that the lengths of the different runs are not equal — the curves are superimposed onto each other; the maximum length of a run is 35 minutes. Right: projection of the data onto 2 dimensions, as determined by the first 2 principal components; for visual purposes, only a (random) subset of the data is plotted.

the most. Note that the variation between the different conditions is not as big as in the simulation. Added to the already large amount of noise existing in the sensors is the fact that the arena is relatively small, and therefore it is harder to separate 'random' behaviour from a more 'controlled' behaviour.

In the low dimensional representation of the data (right of Figure 4), we see the formation of clusters representing the two sides of the wall, and we suspect that due to the small size of the arena we cannot see clearly a cluster representation for when the robot is not near any wall. Note, however, that this 2 dimensional view of the data only accounts for about 64% (in the best case) of the variance. That is, it does not give a full picture of what is going on, and there could be another cluster, but in a different dimension.

3 Discussion

The aim of the analysis in the previous section was to examine the characteristics of the perceptual data that the robot is exposed to under different conditions. We wanted to see if interacting with a robot improves, or enhances, the structure of the data.

Obviously when the robot is unskilled and wandering around the environment randomly on its own, it cannot structure information appropriately for the task, since it has no notion of the task. We can compare this to the situation when the task-knowledge is built-in (hand-crafted), and the robot is then more selective about which parts of the environment it visits. Of course a hand-crafted behaviour is not necessarily the best one, and achieving a good behaviour is hard and sometimes tedious.

A way to avoid trial-and-error type behaviour tuning, is to have the robot interact with an experienced demonstrator, in the form of another robot or a human. As we expected, when the robot followed its demonstrator, or was explicitly guided through the environment with the joystick, the data it was exposed to were better structured than with the random behaviour, and as good as (or better than) the hand-crafted behaviour.

Finally, we have seen how information in the real world is much harder to analyse than in simulation. It is important for a robot to have a good attention system for analysing perceptual information, so that this information can be used for higher cognitive mechanisms, such as learning.

4 Future Implementation of an Attention System

Several unsupervised learning approaches, of which PCA is one, would be suitable as attention mechanisms with which a robot can analyse perceptual information meaningfully. A computational implementation of an attention system such as the one discussed at the beginning of this paper, would need to reflect the individual's current memory 'state' in terms of perceptual structures, how perceptual information activates stored structures for comparisons, and how structures are added and deleted. Furthermore, for a robot interacting with the real world, the system would need to operate on-line.

For these reasons, we believe a Self Organising Feature Map (SOFM) is an appropriate tool, and it has been proven to be very useful in robotic implementation in the past. Further, we are interested in a variation of the SOFM, inspired by Fritzke's Growing Neural Gas algorithm (Fritzke, 1995), where structures grow from experience as required, rather than being specified a-priori (Marsland et al., 2001). We are currently developing such a system that reflects our ideas of attention.

Acknowledgements

Many thanks to Tetsushi Oka (`tetsushi@dai.ed.ac.uk`) for implementations of low-level behaviours on the B21 robot.

References

Breazeal, C. and Velásquez, J. (1998). Toward teaching a robot 'infant' using emotive communication acts. In *Proceedings of 1998 Simulation of Adaptive Behavior (SAB98) Workshop on Socially Situated Intelligence*, Zurich, Switzerland.

Conte, R. (2000). Intelligent social learning. In *AISB'00 Symposium on Starting from Society — The Application of Social Analogies to Computational Systems*, pages 1–11, Birmingham, April 2000.

Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information processing system. *Psychological Bulletin*, 104:163–191.

Fritzke, B. (1995). A growing neural gas network learns topologies. In Tasauro, G., Tourestzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems*, volume 7, pages 625–632. MIT Press, Cambridge MA.

Marsland, S., Nehmzow, U., and Shapiro, J. (2001). Novelty detection in large environments. In *Proceedings of Towards Intelligent Mobile Robots (TIMR) 2001*, Manchester.

Näätänen, R. (1992). *Attention and Brain Function.* Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Posner, M., Snyder, C., and Davidson, B. (1980). Attention and the detection of signals. *Journal of Experimental Psychology General*, 109:160–174.

Scassellati, B. (1998). Imitation and mechanisms of joint attention: A developmental structure for building social skills on a humanoid robot. In *Autonomous Agents '98*.

Sokolov, E. (1963). A neuronal model of the stimulus and the orienting reflex. In Sokolov, E., editor, *Perception and the Conditioned Reflex*, pages 282–294. Pergamon Press.

Spence, K. W. (1937). Experimental studies of learning and higher mental processes in infra-human primates. *Psychological Bulletin*, 34:806–850.

Whiten, A. (2000). Primate culture and social learning. *Cognitive Science special issue on Primate Cognition*, 24. In press.

Zentall, T. R. and Galef Jr., B. G. (1988). *Social Learning: Psychological and Biological Perspectives.* Lawrence Erlbaum Associates, Hillsdale, N.J.