Maintaining Attentional Capacity in a Social Robot

Yuval Marom and Gillian Hayes

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

Abstract

Attention as the perception of change, or event detection, is important for an agent interacting with its physical and social environments. Internal modifications of the controller, in terms of adaptation of a decision threshold used to detect changes, is used to control the level of detail attended to, or the attentional effort. By maintaining effort within some capacity bounds, the agent maintains an attentional drive, which results in not only a 'comfortable' level of processing, but also in desired performance. We present results from simulations of robotic learning by imitation, where an agent learns a task by following a capable teacher agent around the environment.

1 Introduction

One of the problems that we face when we interact with the physical and social world is determining the level of detail to which we should attend, of what we perceive. Put differently, this is a problem of deciding to what extent perceived stimuli need to be significant to justify their going through our mental processing system, and perhaps resulting in learning. Of course, this 'need' is a subjective measure that varies amongst individuals, and depends on current factors, both external and internal.

We are interested in attention as a set of mechanisms that facilitate the *perception of change* [Rensink *et al.*, 1997], or, put it in robotics/agents terminology, *event detection*. We use a *capacity model of attention* [Kahneman, 1973] as a means to control an internal state that mediates decision behaviour. Such a state can be regarded as a *drive* to pay attention, but not too much (or too little).

We do not model emotions directly, but we present several emotion-based autonomy control approaches that are similar to ours.

2 Background

2.1 Perception of Change

Rensink et al. [1997] have performed psychological experiments to test when people perceive changes in images and what influences these perceptions. They conclude that subjects' attention is either *pulled* by transient motion, due to a stimulus' potential merit (saliency), or *pushed* by volitional control, due to internal high-level interests independent of saliency.

We will use statistical significance testing to model saliency. This has been claimed to be an intuitive human cognitive process in the psychophysical literature [Gigerenzer and Murray, 1987; Green and Swets, 1966]. Gadanho and Hallam [1998] follow such an approach, as opposed to most robotic implementations where a fixed, hand-crafted element is used to detect changes (for example [Hayes and Demiris, 1994; Billard and Hayes, 1997; Nehmzow and McGonigle, 1994]). Further, they claim that any significant change in the environment is liable to be captured by a change in an emotional state.

The contribution by Gadanho and Hallam [1998] is that a change is calculated in terms of the observed data, and not some externally designed threshold. However, the desired *significance* of the change is predetermined and fixed. In other words, the significance level of the statistical test is fixed.

We have extended this idea by equipping our robot with a simple strategy for adapting its measure of significance in order to maintain a desired level of attentional effort (see below). In this way we have given the robot the ability to self modify its controller, rather than rely on fine-tuning by the designer.

This leads to the second type of influence mentioned with regards to the experiments by Rensink et al. [1997] — internal control. To model this we use attentional capacity, which we discuss next.

2.2 Attentional Capacity and Effort

Kahneman [1973] offers a capacity model of attention, which places a limit on a human's capacity to perform mental work, but allows this capacity to be freely distributed among concurrent activities according to some allocation policy. The model requires two pieces of information: a stimulus, and a measure of 'effort'. The momentary attentional capacity is determined by the allocation policy.

In this paper we do not deal with multiple concurrent activities, but rather with a single activity. We use a simple allocation policy that places an upper and lower limit on the attentional effort. Effort is maintained within these bounds by attending to more, or fewer, changes in stimuli. The details of the implementation are given in Section 3.

Compared to previous, related work, the modelling of attentional effort as an internal 'monitory' variable is similar to maintaining a 'drive' [Velásquez, 1998; Breazeal, 1998], which is a motivational factor representing an urge.

3 Implementation

In this section we briefly describe the architecture (Figure 1) that is used to model a learner robot, that follows a capable teacher robot around an environment, with the aim of learning a task (such as phototaxis, obstacle avoidance, or maze following). We consider this simple following behaviour to be a form of learning by imitation, similarly to earlier work on robotic learning by imitation [Hayes and Demiris, 1994; Demiris and Hayes, 1996; Billard and Dautenhahn, 1997; Billard and Hayes, 1999]. For more details of this implementation, see [Marom and Hayes, 2000].

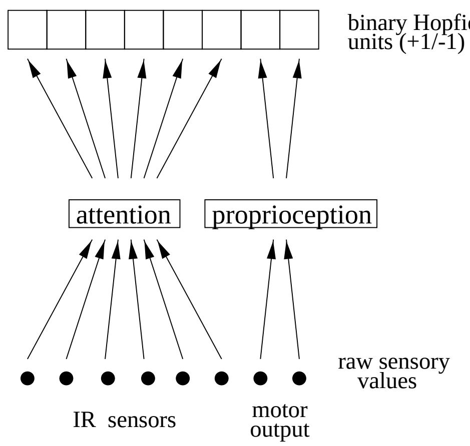

Figure 1: The architecture we have implemented. The attention module processes information from the IR sensors, and the proprioception module processes information from the motor system. The outputs from both modules make up the units of a learning pattern.

The attention module, on the left of Figure 1, takes raw sensor values, and turns them into binary units corresponding to whether a change has been detected or not. This is done as follows, for each sensor:

1. a short-term memory window, of (30) previous values, is created;
2. an average over this window is calculated;
3. this average is compared with the average calculated at the previous time step, using a statistical test;
4. if the test returns a significant result, a change is considered to have occurred, and the appropriate unit, at the top of Figure 1, is turned on; otherwise it is off.

The proprioception module, on the right of Figure 1, inspects the motor values to determine whether a left turn, right turn, or none, has occurred, and sets two binary units (one for left turn, one for right turn) appropriately.

The binary units from both modules make up an associative learning pattern that is fed into a Hopfield associative neural network.

3.1 Learning

In order for learning to take place, the robot must be 'attentive', which occurs when at least one of the units outputed by the attention module is on; further, the robot must be 'socially stimulated' by the teacher, which is achieved through a change in speed.

After training is complete, the robot learner is placed in the environment on its own, and is expected to recall the patterns stored during learning. Since no social cues are now available, the learner can only depend on its internal attention mechanism, which is exactly the same as before. The first part of the pattern is therefore determined by the attention module as before and is fixed; the rest of the pattern is recalled by letting the values of the action units be predicted by the network weights. These units are then translated into actual motor commands that drive the robot.

3.2 Maintaining Attentional Capacity

Within the short-term memory window, the agent keeps track of how frequently it is 'attentive'. We use this frequency as a measure of effort. Using preliminary investigations (Section 4.1), we set the agent's attentional capacity to be within two bounds (50% and 90%). If the momentary frequency is outside these bounds, the threshold for significance is modified accordingly, thus a desired level of effort is maintained. This simple strategy serves to modify the attention 'filter' to control the amount of information attended to (Section 4.2).

Returning to the comparisons with *drives*, Velásquez [1998] defines a drive 'releaser' as a control system that maintains a controlled variable within a certain range. Using his terminology, our drive releaser is natural, or innate, rather than learned.

4 Experiments

Our experiments so far have mainly concentrated on simulations of the Khepera mobile robot [Michel, 1996]. The simulated environment is shown in Figure 2.

The task in which we are interested in these experiments is photo-taxis: finding and approaching light sources. The teacher has this behaviour built-in; it wonders around randomly in the environment until it detects light, it then turns towards it, and speeds up in order to stimulate the learner; once close enough to the light source, it slows down and starts again.

The learner follows behind by using its built-in following behaviour. It processes (ambient) light intensities through its front IR sensors and proprioception

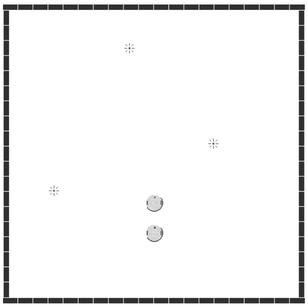

Figure 2: The simulated environment used in our experiments. It consists of a teacher robot followed by a learner robot; and 3 light sources

through its motors, as described in the previous section (Figure 1). It detects a change in speed by referring to its own speed, which varies according to how fast the teacher is moving, since the learner maintains a constant distance through its following behaviour.

4.1 Effort and Performance

We have conducted preliminary systematic experiments, of varying significance thresholds[1] [Marom and Hayes, 2000] (subm.) Each run consisted of 10000 steps of the learning stage, followed by 10000 steps of the recall stage. In the preliminary experiments, 10 runs were carried out for each threshold being tested.

The main conclusion we have drawn from these experiments is that a 'comfortable' region — between 0.2 and 0.4 — in the threshold space exists, where it seems that not too many nor too few patterns are attended to (see Figure 3); Furthermore, this region seems to be a desired one in terms of learning 'quality' and performance.

In order to arrive at these conclusions we have used several evaluation criteria to compare the results. We will discuss two of these criteria in this paper, one related to attentional effort, and the other to performance.

The number of patterns attended to and processed by the learning system, shown in Figure 3, determines the attentional effort exerted by the agent. Note that we are not looking at unique patterns — every time a pattern is attended to, we increment a counter. Different values of the significance threshold will result in different levels of selectivity. That is, the agent will attend to more, or less, patterns.

To test the learned performance, we measured the energy acquired during recall stage. This energy is a function of the light intensity the agent is subject to as it wonders around the environment. Two baselines are used for comparisons: energy acquired as a result of a random behaviour (lower baseline), and a hand-crafted photo-taxis behaviour (upper baseline). The baselines were computed using 50 runs of 10000 steps

[1]We have also run the experiments without the condition that the learner must be socially stimulated by the teacher, and concluded that using social cues is beneficial. See paper for details.

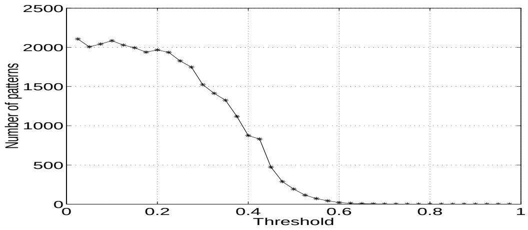

Figure 3: Number of patterns attended to and processed by the learning system, for different significance thresholds. A point on the graphs represent an average over 10 runs.

each, with the appropriate behaviour. The results are shown in the top part of Table 1. We see that the highest energy is achieved in the 'comfortable' region of attentional effort (see bold figure in top part of Table 1).

4.2 Attention Strategy

Using the region of the threshold space found in the preliminary results, we have equipped our agent with a simple strategy to maintain attentional effort — in terms of the number of patterns attended to and learned — such that a significance threshold is self-controlled and modified to lie in this region (see Section 3.2).

Now equipped with this strategy, a run was repeated 50 times, but instead of manually modifying the threshold, we let the agent autonomously adapt it to suit its (built-in) attentional capacity. The evaluations of the performance are shown in middle part of Table 1. We see that on average (bold figure in middle part of Table 1), the agent performs almost as good as the hand-crafted behaviour (bold figure in bottom part of Table 1).

5 Discussion and Conclusion

We have seen that using a simple attention allocation strategy, an agent is able to internally control the level of significance that *it* requires for paying attention to changes in stimulation.

It is important to note that the internal control we have implemented so far is only in terms of 'effort', that is, the load involved in attending to information. Yet this strategy alone is able to produce performance which is not only 'comfortable', but also desirable. We have not even addressed yet the issues of feedback and reinforcement that are available both in the agent's physical and social environments [Mataric, 1994]. Also, one could model emotions in relation to social interactions, involving the perception by others of the self, or valuations of motives of others [Nehaniv, 1998]. We believe that providing our agent with such values would further enhance performance, and we leave this to future work.

Of course, we have only experimented with a single task, and it would be interesting to extend our model to deal with multiple tasks, using multiple sensor modalities. This might require the kind of external feed-

energy ($\times 10^7$)	mean	std. dev.	max
th < 0.2	2.400	0.020	2.430
$0.2 \leq$ th ≤ 0.4	2.474	0.078	**2.594**
th > 0.4	2.337	0.016	2.366
attention strategy	**2.800**	0.094	3.012
random behaviour	2.328	0.056	2.427
hand-crafted behaviour	**2.905**	0.080	3.089

Table 1: Energy acquired during recall stage is used as performance evaluation criteria in preliminary experiments of various significance thresholds (denoted by 'th'), and consequent experiments implementing an attentional strategy that maintains attentional effort according to capacity. Energies acquired as a result of a random behaviour, and a hand-crafted photo-taxis behaviour, can be used as lower and upper baselines, respectively.

back we have just mentioned, but also perhaps some more internal motivational factors other than effort, or attentional 'drive', and might entail a more detailed and direct analysis of emotions.

We did not discuss the topic of social learning in this paper, but we hope that the attention system we have designed would help us analyse in more detail issues concerning social learning, and model them on artificial systems.

References

[Billard and Dautenhahn, 1997] Aude Billard and Kerstin Dautenhahn. Grounding communication in situated, social robots. In *Proceedings of Towards Intelligent Mobile Robots Conference*, 1997.

[Billard and Hayes, 1997] Aude Billard and Gillian Hayes. Transmitting communication skills through imitation in autonomous robots. In *Proceedings of the Sixth European Workshop on Learning Robots*, 1997.

[Billard and Hayes, 1999] Aude Billard and Gillian Hayes. DRAMA, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behaviour Journal*, 7(1), 1999.

[Breazeal, 1998] Cynthia Breazeal, (Ferrell). A motivational system for regulating human-robot interaction. In *AAAI98*. 1998.

[Demiris and Hayes, 1996] John Demiris and Gillian Hayes. Imitative learning mechanisms in robots and humans. In *Proceedings of the 5th European Workshop on Learning Robots*, pages 9—16, 1996.

[Gadanho and Hallam, 1998] Sandra Clara Gadanho and John Hallam. Emotion-triggered learning for autonomous robots. In *SAB'98 Workshop on Grounding Emotions in Adaptive Systems*, 1998.

[Gigerenzer and Murray, 1987] Gerd Gigerenzer and D.J. Murray. *Cognition as Intuitive Statistics*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1987.

[Green and Swets, 1966] D. M. Green and J. A. Swets. *Signal Detection Theory and Psychophysics*. Wiley, New York, 1966.

[Hayes and Demiris, 1994] G. M. Hayes and J. Demiris. A robot controller using learning by imitation. In *Proceedings of the 2nd International Symposium on Intelligent Robotic Systems*, pages 198—204, 1994.

[Kahneman, 1973] D. Kahneman. *Attention and Effort*. Prentice Hall, Englewood-Cliffs, New Jersey, 1973.

[Marom and Hayes, 2000] Yuval Marom and Gillian Hayes. Perception of change and social cues in robotic imitation. In *Proceedings of Autonomous Agents 2000*, Barcelona, June 2000. To appear.

[Mataric, 1994] Maja J Mataric. Learning to behave socially. In *Proceedings, From Animals to Animats 3, Third International Conference on Simulation of Adaptive Behavior*, pages 453–462, 1994.

[Michel, 1996] Olivier Michel. Khepera simulator version 2.0, 1996. Robot simulator written at the Univeristy of Nice Sophia-Antipolis. Downloadable from the World Wide Web at `http://diwww.epfl.ch/lami/team/michel/~khep-sim`.

[Nehaniv, 1998] Chrystopher Nehaniv. The first, second and third person emotions: Grounding adaptation in a biological and social world. In *SAB'98 Workshop on Grounding Emotions in Adaptive Systems*, pages 43–47, August 1998.

[Nehmzow and McGonigle, 1994] Ulrich Nehmzow and Brendan McGonigle. Achieving rapid adaptations in robots by means of external tuition. *From Animals to Animats*, 3, 1994.

[Rensink *et al.*, 1997] R. Rensink, J. O'Regan, and J. Clark. To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8:368—373, 1997.

[Velásquez, 1998] J. Velásquez. Modelling emotion-based decision-making. In *Proceedings of the 1998 AAAI Fall Symposium Emotional and Intelligent: The Tangled Knot of Cognition*, Orlando, FL, 1998. AAAI Press.