Division of Informatics, University of Edinburgh

Institute of Perception, Action and Behaviour

Preliminary approaches to attention for social learning

by

Yuval Marom, Gillian Hayes

Informatics Research Report EDI-INF-RR-0084

Division of Informatics **July 1999**
http://www.informatics.ed.ac.uk/

Preliminary approaches to attention for social learning

Yuval Marom, Gillian Hayes

Informatics Research Report EDI-INF-RR-0084

DIVISION *of* INFORMATICS
Institute of Perception, Action and Behaviour

July 1999

In Proceedings of ACAI'99 Workshop on Biologically Inspired Machine Learning, pp. 8–18, Crete

Abstract :
In social learning, an individual benefits from interacting with its social environment, to acquire new competencies and skills. Social enhancement is a term used to classify all social influences on an individual's performance. An example of this is stimulus enhancement, where one or more individuals, present in a learner's environment, influence the learner's probability of exposure to one set of stimuli rather than others. The learner can take advantage of these enhancements to further reduce the amount of input it has to deal with, by paying attention. Attention is a collection of mechanisms that determine the significance of stimuli. We argue that attention and stimulus enhancement can be used as tools for learning, and discuss their individual and mutual contributions to learning. We present two preliminary, statistical approaches to the modelling of attention, and point out several issues, problems, and possibilities for further work that arise.

Keywords :

Preliminary Approaches To Attention For Social Learning

Yuval Marom and Gillian Hayes

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

ABSTRACT

In social learning, an individual benefits from interacting with its social environment, to acquire new competencies and skills. Social enhancement is a term used to classify all social influences on an individual's performance. An example of this is stimulus enhancement, where one or more individuals, present in a learner's environment, influence the learner's probability of exposure to one set of stimuli rather than others. The learner can take advantage of these enhancements to further reduce the amount of input it has to deal with, by paying attention. Attention is a collection of mechanisms that determine the significance of stimuli. We argue that attention and stimulus enhancement can be used as tools for learning, and discuss their individual and mutual contributions to learning. We present two preliminary, statistical approaches to the modelling of attention, and point out several issues, problems, and possibilities for further work that arise.

INTRODUCTION

In the study of intelligence, the social animate dynamics are just as important as the inanimate environmental dynamics (Dautenhahn, 1995). The *social intelligence hypothesis* states that as well as dealing with the physical world, intelligence also concerns dealing with other individuals, whereby we use each other as 'social tools' (Dautenhahn, 1995).

When we think about an individual 'using' another to acquire knowledge, experience, skills, etc. we are almost immediately dealing with the concept of learning by imitation, or observation. It seems intuitive to take advantage of an individual's expertise by mimicking it and trying to understand why a particular response was useful in a particular situation. The reason it seems intuitive is because this is what we humans do, from a very early developmental stage (Galef, 1988; Dautenhahn, 1995), and even from birth (Meltzoff and Moore, 1983).

An autonomous robot can be modelled as an infant, where an intuitive desire and ability to learn can be assumed, together with built-in skills waiting to be triggered and receive the right training. Social learning is possible when a more able individual is present. In human-robot-interaction scenarios, a caretaker-infant dyad can be employed in the modelling (Ferrell, 1998; Ferrell and Scassellati, 1998; Scassellati, 1998). Robot-robot interactions have also been useful in the study of artificial social intelligence (Hayes and Demiris, 1994; Billard and Hayes, 1999; Dautenhahn, 1995).

Learning implies a continuous processing of information, and even humans do not and *cannot* process all the information available to them via the perception system. Hence we need mechanisms for deciding what we deem to be important; what deserves our attention.

The purpose of the work reported here is to devise pre-learning attentional mechanisms, which would reduce the amount of information that an individual has to deal with, and hence enhance learning. This paper presents some initial statistical approaches we have taken to model attention. The question of how plausible these statistical tools are from a biological and psychological point of view remains unanswered, although we show that computationally they are effective. In this paper we do not deal with any actual learning per se, only with what precedes it. In future work we intend to investigate how our mechanisms affect learning, if at all.

In our approach, we conceptualise attention, and another concept called *social enhancement* (see below), as *tools* for learning. In the following two sections, respectively, we present some background material for these two concepts, and discuss how they are used as tools for learning.

BACKGROUND

Attention

Attention is a collection of mechanisms that determine the significance of stimuli (Kahneman, 1973). We are interested in the kind of attention that facilitates perception of change. According to psychological experiments performed by Rensink et al. (1997), perception of change is possible with the use of an attentional bottleneck, where attention is attracted to various parts of the environment based on high-level interests.

In these experiments, Rensink et al. (1997) have introduced a distinction between *central* and *marginal* areas of interest, and shown that subjects exhibit different attentional behaviors between these areas.

Social Enhancement

In social learning, an individual benefits from interacting with its social environment, to acquire new competencies and skills. In other words, the existence of one or more other individuals in its perceived environment aids an individual to learn as it negotiates unknown environments. Psychologists use the term *so-*

cial enhancement to refer to all social influences on an individual's performance (Galef, 1988; Hogan, 1988).

One form of such influence is *local enhancement* (Galef, 1988) or *stimulus enhancement* (Spence, 1937), where one (or more) individual (the teacher, or demonstrator) actively manipulates the perceived environment of another (the learner). The purpose of these manipulations is to direct the attention of the learner to the relevant stimuli of the task to be learned.

Bennet Galef argues that local/stimulus enhancement can occur from a tendency of the individual to approach conspecifics, and from alterations conspecifics have made in the environment or objects they have contacted (Galef, 1988). According to Galef, these situations lead to a change in the probability of exposure to one set of stimuli rather than others. The mechanisms we present in this paper are based on the first of these scenarios.

CONTRIBUTIONS TO LEARNING

As stated earlier, we conceptualise attention and social enhancement as *tools* for learning. In this section we discuss how these tools contribute to learning *and* to each other, from a modelling point of view. Both the learner and the teacher take an active role in this process, shown in Figure 1.

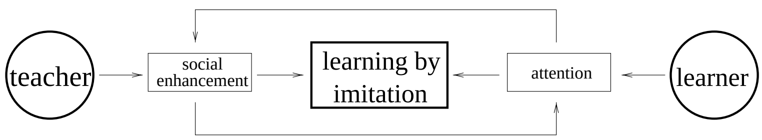

Figure 1: Tools for learning by imitation

Social enhancement, in the form of stimulus enhancement, contributes to learning because it exposes the learner to the relevant stimuli, a subset of the total stimuli, of the task to be learned. The issue of relevance is determined by the teacher. Attention contributes to learning because it filters the information required for processing by the learning system, according to the significance of changes. Here it is the learner who determines when something is significant and worth further processing. Together these tools reduce the amount of information storage required and hence increase learning speed.

Stimulus enhancement simplifies the modelling of attention because the learner can be confident that the teacher will only guide it through useful stimuli. The attention system, therefore, only needs to decide *when* something is important (by deciding that it is significant) and not *if* it is important. In a sense, the desire to be 'socially enhanced' is a kind of high level interest that drives the attention (see Background section). This kind of facilitation by the teacher is evident in infant-caretaker relationships and is often termed *scaffolding* (Ferrell and Scassellati, 1998; Wood et al., 1976).

The learner needs to decide when it is being 'socially enhanced', in other words when is the teacher trying to get its attention. It is the job of the attention system to make these decisions. So attention contributes to social enhancement.

MECHANISMS OF ATTENTION

As outlined in the Background section, the learner maintains close contact with the teacher, and is consequently guided through the task. It can be confident that it will benefit from staying with the teacher, due to social enhancement. Therefore by adapting an imitation strategy (in our case *following*), all it needs to decide is *when* something is different, that is when it is perceiving a significant change.

In the preliminary stages of our work, we have identified two statistical mechanisms for modelling attention, or the perception of change. The first of these uses short-term memory to compare new information with immediately preceding experience. Experience is computationally modelled as the average of the perceived sensory inputs, over a fixed short-term memory window. When this comparison yields a significant difference, attention is turned on, otherwise it is off.

The second mechanism utilises known knowledge about the sensory noise produced by the simulator to calculate the parameters of a "base-line" statistical distribution. This fixed base-line distribution represents the stimuli under "normal" conditions, that is when no stimulation is expected. New experience is calculated as before, but is now compared with the experience under normal conditions. This method therefore uses long-term as well as short-term memory. In a real, physical world one would not have a-priori knowledge about the noise in the sensors, and would therefore have to estimate it using initial experience. We intend to implement this in future work.

Perceptions are grouped into three "regions of interest": *central*, *marginal*, and *peripheral*. A separate memory is allocated to each region, with a different window size for each. The central region has the largest-sized window, followed by the marginal region, and the peripheral region. This kind of memory decomposition is favourable for two reasons:

1. it reduces the amount of computation and storage required to process the input, and
2. it follows findings, found by Rensink et al. (1997) in their experiments, which state that subjects identified areas of interest, where more "interesting" observations were kept longer in memory, and less "interesting" ones were overwritten more frequently (Rensink et al., 1997).

This could also be thought of as modelling some kind of high-level motivation that is used to govern attention, although the number of regions and their corresponding window sizes are fixed. A more flexible, dynamic decision process would be a better model of such high level mental activities, and justifies future work in the area.

EXPERIMENTS

Two experiments, corresponding to the two methods described in the previous section, were carried out using simulations. A Khepera mobile robot simulator Michel (1996) was used in these experiments. A diagram of the Khepera robot (both real and simulated) is shown in Figure 2. The robot has two motors and eight infra-red sensors, capable of detecting distance and light.

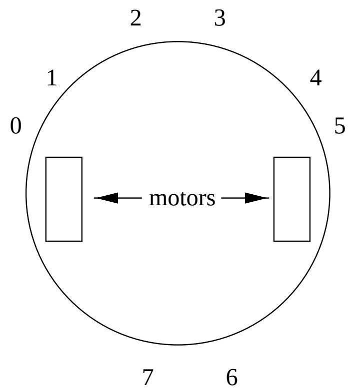

Figure 2: A diagram of the Khepera robot. It has two motors, and 8 IR sensors labelled 0 to 7.

In the experiments reported here two robots are used – a teacher and a learner. The teacher exhibits a photo-taxis behavior. It is instructed to speed up when it is approaching a light source; stop when it gets there; escape from the light source for a given amount of time; and start again. The learner exhibits a follow-teacher behavior[1]. A diagram of the simulated environment is shown in Figure 3.

The learner's light sensor values are grouped, as discussed in the previous section, into three regions as follows: central region – sensors 2 and 3, marginal region – sensors 1 and 4, and peripheral region – sensors 0 and 5; sensors 6 and 7 are not being used (see Figure 2). The window sizes for the different regions are determined empirically.

The learner will eventually learn the task of approaching light sources. During imitation, it needs to decide when it is seeing something useful, and hence attend to it. Note that since the action *approach-light* is not instantaneous but rather falls within an (finite) interval, the learner also needs to decide for how long to maintain attention.

Attention to Local Changes

In the first experiment, successive memories were compared using averages. These differences were regarded as the stimuli, and a separate stimulus was computed for each region of interest. The total stimulus was computed as a weighted sum of the different regions, with the central region having most

[1]Both teacher and learner are also equipped with an avoid-obstacle behavior. However, since the learner is always directly behind the teacher, it rarely utilises it.

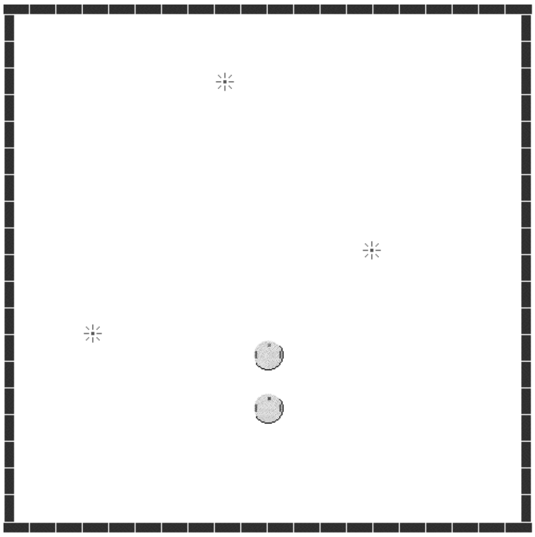

Figure 3: A diagram of the simulated environment, with three light sources and two robots – a teacher, followed by a learner (the dot on a robot represents its front).

weight, followed by the marginal and peripheral regions. This total stimulus was put through a threshold to determine if attention should be turned on. This threshold was determined empirically. Figure 4 presents a plot of the perceived environment (light intensities) and total stimulus, associated with three encounters with a light source.

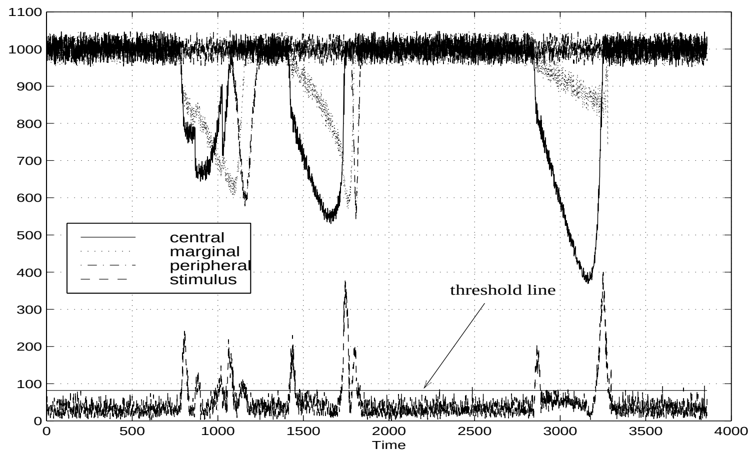

Figure 4: **Local attention**. Perceived environment versus total stimulus. The grouped light intensity values appear at the top of the plot (small values indicate strong intensities); the total stimulus and the threshold for attention (scaled to fit on the plot) appear at the bottom of the plot. This particular plot represents three encounters with a light source.

We notice from Figure 4 that each encounter with a light source results in two attentional peaks: at the beginning of the *approach-light* action and at its termination. We expect this behavior since we are using local attention here:

once a light source is found, there's a big local change in stimulus, but as the robot gets closer to the light, the immediate changes are not very big, until the robot passes the light and the intensities drop to zero, resulting in a big local change.

Attention to Global Changes

In the second experiment, the parameters of a baseline (uniform) distribution were calculated, using known information about the random noise produced by the simulator[2]. This distribution was treated as the group of values of the perceived environment under "normal" conditions, that is when no light is visible (see Figure 5).

Each window (distribution) of values was compared with the baseline distribution, and a measure of difference was computed (a z-statistic). High values indicated (statistically) significant differences, and in these situations attention was turned on. The threshold used was the 1% significance-level critical value for determining significance (2.33 – from statistical tables).

A separate stimulus (value of a z-statistic) was computed for each region, and the highest stimulus of the three was used as the total stimulus. Consequently, as long as a significant change in stimulus was perceived in any of the areas of interest, attention was turned on.

Figure 5 presents a plot of the perceived environment (light intensities) and total stimulus, associated with three encounters with a light source (note: this is not the same run as the one presented in the first experiment).

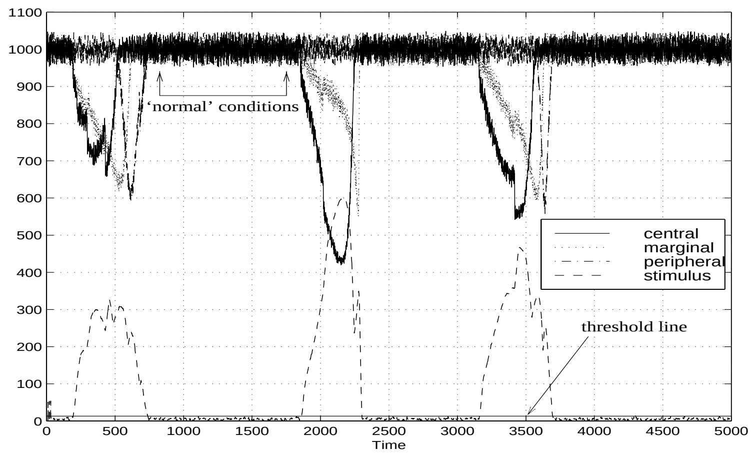

Figure 5: **Global attention**. Perceived environment versus total stimulus. The grouped light intensity values appear at the top of the plot (small values indicate strong intensities); the total stimulus, given as values of a z-statistic, and the 1% significance-value threshold line (scaled to fit on the plot) appear at the bottom of the plot.

---

[2]A uniform distribution, with values in the range $[a, b]$ has mean $(a+b)/2$ and variance $(b-a)^2/12$.

In contrast to the first experiment (see Figure 4), we see from Figure 5 that attention is turned on for the duration of the *approach-light* action, since the perceived environment is significantly different during the action than it is at other times. Attention is being applied globally.

Evaluation of Attention

In order that learning is triggered at the right times, we need to examine whether attention is working correctly. We see from Figures 4 and 5 that in terms of the learner's own perceptions, it is paying attention at the right times. However, we need to verify that this is indeed the case, using a more informed source: the teacher.

The teacher signals that it is doing something important (approaching the light) by speeding up, so we can compare the times when it does so with the learner's attentional behaviour. Figures 6 and 7 present plots of the teacher's speed, superimposed on the learner's stimuli, in the first and second experiments, respectively.

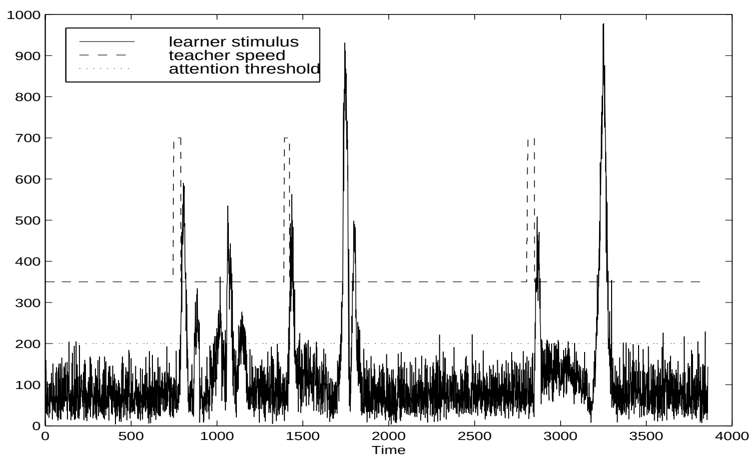

Figure 6: Comparison of teacher's speed (scaled) and learner's stimuli with local attention.

We can see from these plots that, allowing for a time-lag due to the distance between teacher and learner, attention is turned on at the right places. However, we also see that the teacher slows down quite a long time before the learner stops paying attention. This is due to the fact that when the teacher gets too close to the light (according to a threshold), it stops and escapes for a certain time at normal speed, even though it might still be close to the light. The learner is also still close to the light, but the only information it has at this point is that the light intensities haven't changed much in value, so it is maintains its attention.

The purpose of the teacher when speeding up is to attract the attention of the learner, because this signals that the teacher is doing something important. Currently, only exteroceptions (the perceived changes in the environment) affect the attention decision process. It would be beneficial for the learner to incorpor-

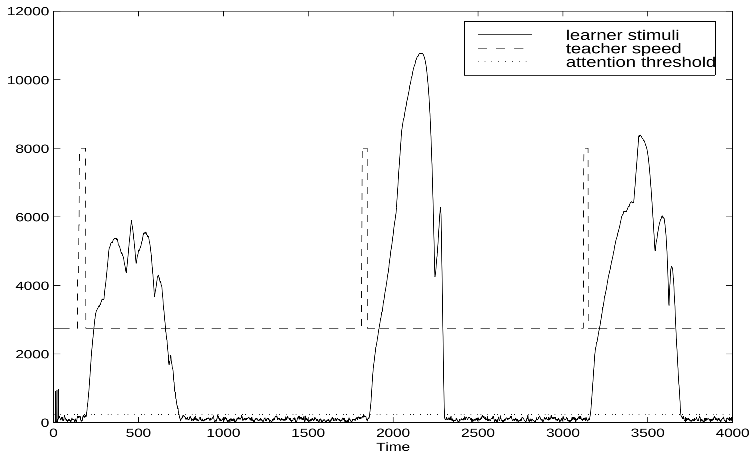

Figure 7: Comparison of teacher's speed (scaled) and learner's stimuli with global attention.

ate its own proprioceptions, obtained as a result of imitating the teacher, in this process, and therefore make better use of the stimulus enhancement provided by the teacher (see 'Social Enhancement' in Background section). In future experiments we intend to implement this capability.

DISCUSSION

There are two main differences between the two experiments reported in the previous section, concerning the following issues:

1. how a change in stimulus is calculated – local versus global;
2. how the significance of the change, i.e. the threshold for attention, is determined.

Local Versus Global Attention

In the first experiment, the learner uses local changes of stimulus to determine its attentive state, whereas in the second one, it uses global changes. These two approaches provide the learner with different information, and are useful for different reasons. The first approach gives the learner the start and finish points of the attention interval, whereas the second gives it the whole of the interval.

In order to determine which of these approaches is more biologically and psychologically plausible, one needs to go into a more detailed investigation of attention than was carried out prior to the preliminary approaches reported in this paper. It seems possible, however, that a combination of these two mechanisms could provide a useful model. We know that humans use both short- and long-term memory, but to what extent is each of these used to pay attention? In future work, we hope to develop and implement a biologically- and psychologically- inspired model of attention.

Threshold for Attention

In the first experiment, the threshold for attention was determined empirically, based on initial results. This value was modified a number of times to prevent the learner from paying attention “unnecessarily”. This is a hand-crafted approach that might not scale up very well. In the second experiment, the threshold was determined using statistical tables. These tables contain predetermined values, used globally for statistical analysis. In practice, one such look-up table could be made available for the robot to use for determining when it is seeing something important. The key thing to note here is that the data are standardised to be comparable with this fixed distribution of values. So the comparison is possible regardless of the shape and size of the input data, as long as enough data are used (that is, the window sizes are large enough)[3].

Therefore, for an autonomous robot, the second mechanism is much more computationally effective, because the robot can adapt to its environment. In our specific implementation, this is not entirely true, however, because we are using known information from the simulator. However, robust (non-parametric) statistical tools do exist for more realistic situations.

CONCLUSION

We have seen how social enhancement can be used as a tool for social learning. If the learner can intuitively assume the presence of social enhancement (as in the infant-caretaker dyad), this could potentially have important implications to the modelling of attention. The learner need only figure out when something *significant* is happening, because it can be confident that this will be useful for its learning task.

Biological and psychological aspects of attention have not been fully explored yet. However, two computational mechanisms were shown to be useful, in a simulated world. In future work, we plan to extend these mechanisms, and others, to real robots.

Other issues that have not been addressed in this paper are multiple sources of stimulation, and multiple behaviors. In these situations, the learner has to decide how to distribute and allocate attention to deal with these multiple sources, and choose the appropriate behaviors. This might require some kind of internal high level motivational factor, such as arousal. Kahneman (1973), in his book *Attention and Effort*, identifies the notion of attentional *capacity*, which is limited and can be distributed. Furthermore, he analyses the relationship between capacity and arousal.

Our overall research goal is to develop a full model of attention based on experiments performed on humans, such as (Kahneman, 1973; Rensink et al.,

[3]this is known as the Central Limit Theorem.

1997). This would perhaps involve using more biologically-inspired models such as neural networks.

References

Billard, A. and Hayes, G. (1999). DRAMA, a connectionist architecture for control and learning in autonomous robots. *Adaptive Behaviour Journal*, 7(1). (In press).

Dautenhahn, K. (1995). Getting to know each other – artificial social intelligence for autonomous robots. *Robotics and Autonomous Systems*, 16:333—356.

Ferrell, C. (1998). A motivational system for regulating human-robot interaction. In *AAAI98*.

Ferrell, C. and Scassellati, B. (1998). Infant-like social interactions between a robot and a human caretaker. *Special Issue of Adaptive Behavior on Simulation Models of Social Agents*. In submission.

Galef, Jr., B. G. (1988). Imitation in animals: history, definition, and interpretation of data from the psychological laboratory. In *Social Learning: Psychological and Biological Perspectives*, pages 3—28. Lawrence Erlbaum Associates, Hillsdale, N.J.

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

Hogan, D. E. (1988). Learned imitation by pigeons. In *Social Learning: Psychological and Biological Perspectives*, pages 225—238. Lawrence Erlbaum Associates, Hillsdale, N.J.

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

Meltzoff, A. N. and Moore, M. K. (1983). Newborn infants imitate adult facial gestures. *Child Development*, 54:702–709.

Michel, O. (1996). Khepera simulator version 2.0. 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`.

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

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

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

Wood, D., Bruner, J. S., and Ross, G. (1976). The role of tutoring in problem-solving. *Journal of Child Psychology and Psychiatry*, 17:89—100.