A Model of Habituation Applied to Mobile Robots

Stephen Marsland, Ulrich Nehmzow and Jonathan Shapiro
Department of Computer Science
University of Manchester
Manchester M13 9PL
United Kingdom
(smarsland, ulrich, jls)@cs.man.ac.uk

26.3.99

Abstract

We transfer a neuronal model of the biological phenomenon of habituation, where repeated stimuli begin to be ignored, onto a mobile robot. Results demonstrating the validity of the model are given. Applications of habituation to mobile robots are suggested and areas of further research are presented.

1 Introduction

Habituation, the mechanism by which the brain learns to ignore repeated stimuli, is the most basic form of plasticity within the brain. It has been investigated thoroughly in recent years by biologists in a number of different animals, from the sea snail *Aplysia* [1, 2, 3, 5] through to cats [11] and toads [4, 14]. Of itself, habituation is an interesting phenomenon and therefore worthy of investigation, but there are several worthwhile uses to which it can be put. For example, a novelty detector of the type proposed by Kohonen [8] is effectively an habituation mechanism since input patterns which are presented to the network repeatedly are ignored.

Habituation can be thought of as a way of defocusing attention from features which are seen often, allowing the agent (be it an animal or a robot) to concentrate fully on other, newer experiences. This paper transfers a model of habituation suggested in the biological literature onto a mobile robot and shows that the robot's response to repeated stimuli is indeed reduced over time, as predicted by the model. Experiments which parallel those carried out on toads are performed with analogous results.

2 A Brief History of Habituation

Evidence of habituation can be seen most clearly in *Aplysia*. This mollusc has a gill which is withdrawn when its siphon is touched. However, repeated gentle stimulation of the siphon results in an habituated response meaning that the gill is withdrawn less and finally not at all. Repeated series of training show that, while the defensive withdrawal returned over time (dishabituation), further stimulation habituates faster. This demonstrates that habituation is a two stage process, there is a short-term effect lasting for a few minutes and a long-term effect lasting for hours. For a more comprehensive review, see [1, 19].

Although habituation causes a decrement in response, the full response can be regained without allowing the animal to rest, for example by training with a different stimulus. This makes it different to fatigue. For a complete definition of habituation, see Thompson and Spencer [12].

There have been several attempts to model the effects of habituation computationally and to explain the interaction between habituation and dishabituation, the process whereby an habituated response returns to its original strength. In 1970, Groves and Thompson [6] suggested that dishabituation was an instance of sensitization and therefore an independent construct which interacts with habituation to produce a net response. Their model was used by Stanley [10] to simulate data from experiments of

Proc. TIMR 99 "Towards Intelligent Mobile Robots", Bristol 1999.
Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 - 6161.
Report number UMCS-99-3-1. http://www.cs.man.ac.uk/cstechrep/titles99.html

the spinal cord of a cat. He described the decrease of synaptic efficacy y by the first-order differential equation

$$\tau \frac{dy(t)}{dt} = \alpha(y_0 - y(t)) - S(t) \tag{1}$$

where y_0 is the initial value of y, $S(t)$ is the external stimulation and τ is a time constant governing the rate of habituation, while α controls the recovery rate. A graph showing the effects of this equation is given in figure 1.

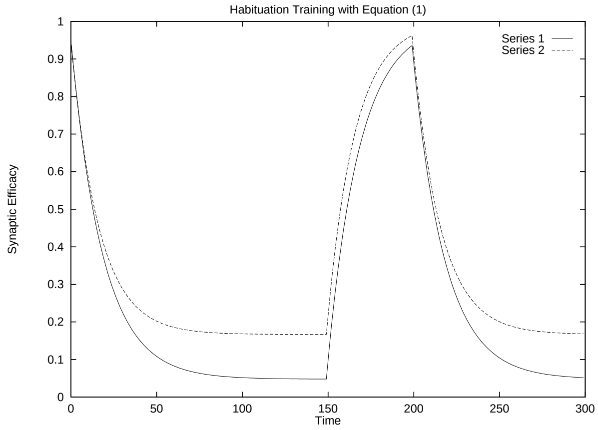

Figure 1: An example of how the synaptic efficacy drops when habituation occurs. This example uses Stanley's model (equation 1). In the first, descending part of the graph, a stimulus $S(t) = 1$ is presented continuously. This changes to $S(t) = 0$ at $t = 150$ where the synaptic efficacy rises again, and becomes $S(t) = 1$ again at $t = 200$, causing another drop. The two curves show different values of the constants, in series 1 $\alpha = 1.05$ and in series 2 $\alpha = 1.2$. In both, $\tau = 20$ and $y_0 = 1.0$.

Unfortunately this model displays only the short-term effects. The first model incorporating long-term habituation was produced by Wang and Hsu [17], who used these equations:

$$\tau \frac{dy(t)}{dt} = \alpha z(t)(y_0 - y(t)) - \beta y(t)S(t), \tag{2}$$

$$\frac{dz(t)}{dt} = \gamma z(t)(z(t) - 1)S(t). \tag{3}$$

This pair of equations give an S-shaped curve which displays both short-term and long-term effects of habituation. β controls the speed of habituation and the input $S(t)$ is gated by being multiplied by $y(t)$ instead of the direct input of equation 1. The new variable, $z(t)$ changes slowly compared with $y(t)$ and is used to control the rate of recovery. It has a single point of inflection, above which recovery is rapid, corresponding to short-term memory, while lower down long-term effects dominate.

In 1978 Ewert and Kehl [4], performing experiments where a number of different prey were shown to toads, showed that while some changes in dummy prey dishabituated others, some did not affect the habituated response. They found that there was a hierarchy of habituation, and that only stimuli further up in the hierarchy could dishabituate previously habituated responses. The most likely explanation for this (see Wang and Ewert [16] for more details) is that dishabituation is a return to normal behaviour and cross-talk between the habituation and dishabituation processes cause the failure of dishabituation.

3 The Biological Experiments

The discovery of the habituation hierarchy mentioned above has been of very great interest to investigators. A huge number of experiments have been conducted in attempts to discover why it exists.

Most notable are those of Ewert and Kehl [4], Lara and Arbib [9], Wang and Ewert [16] and Wang and Arbib [13, 14, 15].

All the experiments were performed on the toad *bufo bufo*, and are concerned with the ability of the toad to recognise different prey dummies and habituate its response to their presentation so that it no longer attempts to catch the prey. Their results show that habituation displays locus specificity, that is it depends where on the retina the prey is seen, and hierarchical stimulus specificity so that a second stimulus on the same retinal location will only get a response if it is higher in the habituation hierarchy. This means that habituating to one stimulus causes all stimuli at the same level in the hierarchy, and those below, to be habituated too.

All the experiments were set-up in the following way. A toad sat in a cylindrical glass case while a variety of prey dummies were mechanically moved across the visual field of the toad. By presenting the shape at the same height and position relative to the toad the image was made to appear on the same part of the retina. If the toad recognised the dummy as prey it would follow its motion by turning to continue watching as the prey moved. Showing the same dummy continuously showed that toads habituate the turning motion if they recognised that the particular prey was unreachable, that is that the animal had previously failed to catch that prey.

The original prey dummy was a rectangle of black cardboard measuring 20mm by 5mm. This moved clockwise about the toad at a constant distance of 7cm from it.

The presentation was then modified in a number of ways, two of which have analogies in our experiments:

- using a different prey dummy (a different shape) and presenting it so that it appeared on the same part of the retina,
- moving the same dummy so that it appeared on a different part of the retina.

4 The Neural Model

In 1991 and 1992 Wang and Arbib produced a series of papers [13, 14, 15] which attempted to produce a neural model to simulate the results of the experiments described in the previous section. They modelled the tectal relay and anterior thalamus (AT), the areas which process the image taken from the retina ([13]) and extended the model ([15]) with a neural mechanism for the medial pallium (MP), the region in which habituation is thought to take place. Their model of the MP is shown in figure 2 and the equations controlling it are described below. We make no attempt to justify the design, leaving that to its designers.

Their conjecture is that the MP consists of a large number of columns arranged vertically with five layers of cells in each column, each layer consisting of n neurons ($n = 20$ on our robot, although figure 2 shows $n = 5$ for simplicity). In this paper we investigate the properties of one of these columns and show that this displays the behaviour observed in the biological experiments.

One neuron in the AT propagates its response strength to all the P1 neurons together. The membrane potential $m_{p1}(i,t)$ - the internal variable of a cell - of the ith P1 cell, $P1_i$ is defined as

$$\frac{dm_{p1}(i,t)}{dt} = -A_{p1}m_{p1}(i,t) + B_{p1}I(t) + \varrho \tag{4}$$

where $A_{p1} = 1$ is a relaxation parameter and $B_{p1} = 1$ a connection weight from the AT input, $I(t)$. The additional term ϱ is a noise term included to compensate for the fact that $I(t)$ is assumed to be constant during a stimulus presentation. The values of A_{p1} and B_{p1} mean that at its maximum, $m_{p1} = I(t)$, so that the neuron simply passes on the input signal, which is chosen according to which stimulus is received. The output of cell $P1_i$, $N_{p1}(i,t)$ is given by:

$$N_{p1}(i,t) = \begin{cases} m_{p1}(i,t) & \text{if } m_{p1}(i,t) > \theta_i \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where θ_i is a threshold which increases monotonically with i and is designed to make cells with larger i harder to fire, so that stronger stimuli (those with larger $I(t)$) have a larger number of neurons firing.

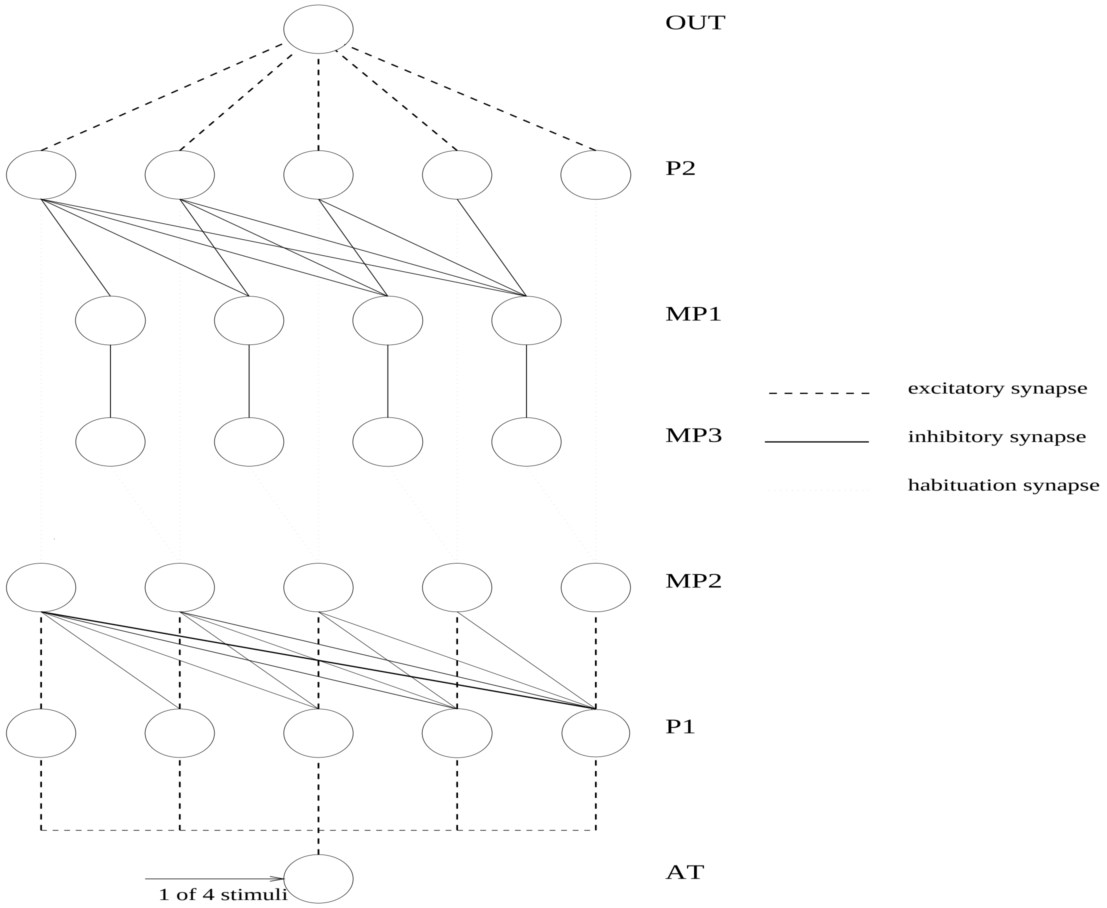

Figure 2: A diagram of the hypothetical neural layout of a column of the medial pallium of a toad, showing the different types of synapse. All the connections are shown. An inhibitory synapse causes a reduction in the activity of the neurons it connects with, while an excitatory synapse increases the activity. The habituation synapses are the only plastic connections and are governed by equations 2 and 3. The thickness of the lines in the inhibitory synapses between P1 and MP2 describes the strength of the connections, the thicker the stronger.

This has implications for the MP2 layer, as can be seen in figure 2, since neurons on the right of P1 have an inhibitory input to all neurons to the left of them in the next layer. The chosen definition of θ_i is: $\theta_i(i) = 0.8i/n + 0.1$. This is only a threshold, once a value has exceeded it the value of θ_i does not affect the actual output of the neuron.

An MP2 neuron receives inhibitory input from P1 neurons to the right of it, with synaptic strengths increasing with the horizontal distance between the two neurons, and an excitatory input from the neuron immediately below it, shown by the following equations:

$$\frac{dm_{mp2}(i,t)}{dt} = -A_{mp2}m_{mp2}(i,t) + (B_{mp2} - m_{mp2}(i,t))I_{ii}(t) + m_{mp2}(i,t)\sum_{j>i} I_{ij}(t) \tag{6}$$

where $I_{ij}(t)$ represents the input from neuron $P1_j$ and is equal to $W_{ij}N_{p1}(j,t)$ for $j \geq i$ with $W_{ii} = 1$ and $W_{ij} = (i-j)/3 < 0$ for $(j > i)$. m_{p2} is the membrane potential of the MP2 neuron and $A_{mp2} = 0.1$, $B_{mp2} = 1.1$. This means that as $I(t)$ increases, the position of the cell with the largest potential moves to the right. Individually each neuron increases to a maximum value and remains there if the input is held constant over time.

The output from the cell $MP2_i$ is formed by

$$N_{mp2}(i,t) = m_{mp2}(i) + h_1 \tag{7}$$

where $h_1 = 0.6$ is the base activity of the MP2 neurons, the output of the neuron if no input is received.

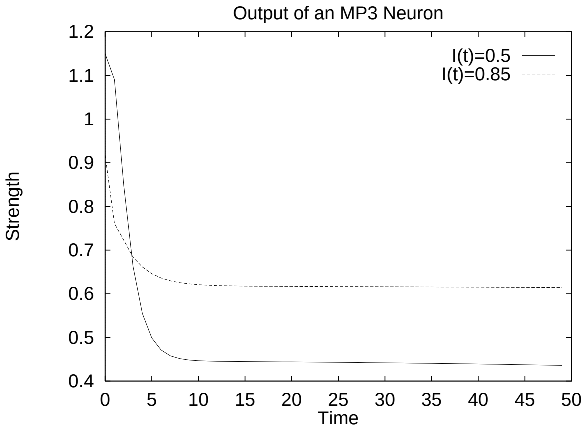

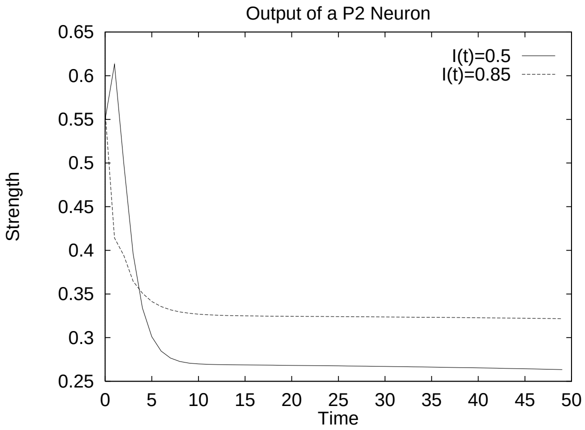

Figure 3: *Left:* The output of an MP3 neuron for different values of input strength with the input held constant over time. The value of the output drops over time for a constant input, with smaller inputs dropping to a lower final value. *Right:* The same effects can be seen in a P2 neuron.

This output is propagated to the MP3 and P2 layers as an excitatory input via the habituation synapses. The output of a typical cell in the MP3 and P2 layers is shown in figure 3 for two different input values. MP3 cells have output $N_{mp3}(i,t)$, the same as their membrane potential, $m_{mp3}(i,t)$ defined by:

$$\frac{dm_{mp3}(i,t)}{dt} = -A_{mp3}m_{mp3}(i,t) + y_i(t)N_{mp2}(i,t) \tag{8}$$

The $y_i(t)$ terms are the habituable weights, whose dynamics are governed by equations 2 and 3 with $S(t) = [N_{mp2} - h1]^+$, where $[x]^+ = \mathrm{Max}(0,x)$ to ensure that the input is at least zero. $A_{mp3} = 1$ and $y_i(0) = 1$. As can be seen on the left of figure 3, if the input is held constant, then the output decreases to some value and then remains there. The stronger the input, the larger the final value. This output forms the sole inhibitory input for the MP1 layer, whose membrane potential, $m_{mp1}(i,t)$ is given by

$$\frac{dm_{mp1}(i,t)}{dt} = -m_{mp1}(i,t) + h_2 - B_{mp1}N_{mp3}(i,t), \tag{9}$$

for $h_2 = 0.6$ and $B_{mp1} = 1$. The output is

$$N_{mp1}(i,t) = [m_{mp1}(i,t)]^+. \tag{10}$$

Finally the P2 layer receives an inhibitory input from the MP2 layer which reduces with further presentations of a stimulus, and further inhibitory input from all neurons to the right of it in the MP1 layer, although in this case all inputs are weighted equally. The outputs of this layer,

$N_{p2}(i,t) = [m_{p2}(i,t)]^+$ with

$$\frac{dm_{p2}(i,t)}{dt} = -A_{p2}m_{p2}(i,t) + y_i(t)(N_{mp2}(i,t) - h_1) - B_{p2}\sum_{j>i}[N_{mp1}(i,t) - C_{mp1}]^+ \tag{11}$$

($A_{p2} = 1$, $B_{p2} = 0.1$) are summed up by the OUT cell. The behaviour of the P2 layer is similar to that of the MP3 layer, as can be seen in figure 3.

C_{mp1} is the resting activity of an MP1 neuron before any habituation occurs. Using equations 7, 8 and 9 we can see that

$$C_{mp1} = h_2 - h_1 * y_0 * B_{mp1}. \tag{12}$$

The values of all constants were the same as those used by Wang and Arbib ([15]), either declared explicitly there or inferred from their text.

5 The Robotic Version

We implemented the neural model shown in figure 2 on an autonomous mobile robot. It was then possible to investigate the properties of the model and to see that the habituation hierarchy exists through the reactions of the robot to different stimuli.

The robot used was a small mobile robot built from the FischerTechnik range of interlocking parts with a Motorola 68HC11 microcontroller, an integrated circuit with a 6811 microprocessor and 32K of RAM. It is programmed using MIT's Interactive C control system. The robot has a two wheel differential drive system and is equipped with four digital bumper sensors, an infra-red transmission system and four light sensors facing in the cardinal directions. Only the light sensors are used in these experiments.

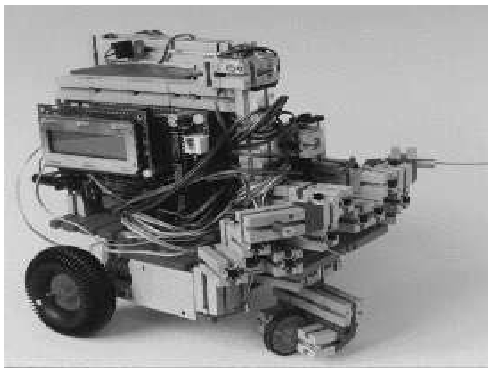

Figure 4: The FischerTechnik mobile robot used in the experiments.

A hard-wired behaviour of phototaxis was introduced as the experimental equivalent of the toads prey-catching behaviour. To imitate different "prey dummies" we varied the light shown. We used four light stimuli, constantly on, flashing quickly, flashing slowly and dim light, labelled S1 to S4 respectively in the results section.

The input to the network (the role of the AT in the toad) was taken by a process which recognised the current lamp state and produced the value predefined for that stimulus. The values chosen for the inputs ($I(t)$ in section 4) were 0.85 for stimulus S1, 0.6 and 0.5 for stimuli S2 and S3 respectively and finally 0.35 for S4. The higher the value, the longer the robot would take to habituate to that stimulus, a feature of habituation. These four values form a hierarchy of habituation, with S1 at the top, S2 and S3 sharing a middle level and S4 at the bottom. For more information, see the discussion in section 8.

The output from the network was used to determine the motor force. The robot was not constrained so that it would always move towards the light, but habituation meant that the speed with which it moved would slow and finally stop as habituation progressed. We could then investigate whether or not showing a different light restarted the motion.

For the purposes of the experiments reported in this paper the light was shone directly onto the light cell at a very short distance using a torch. The intensity of the light played no part in the input to the network, but the light cell with the brightest light impinging on it was the one whose stimulus was fed into the neural model. We therefore needed to ensure that the stimulus we wanted was recognised correctly by the chosen sensor and was the one chosen to be presented to the single column of neurons used in the implementation, which meant ensuring the light was brightest on that sensor.

The experiments were conducted as follows:
A light was shone onto a particular light cell on the robot, which would turn so that it faced the light and then move towards it. However, by moving the light at the same speed as the robot so the light remained on a particular light cell (thus rendering the "prey" unreachable) it was possible to make the robot habituate its response to that so that its speed slowed down gradually and eventually it would stop. The robot did not know that the light was turning with it, but the effect of moving the light with the robot was to keep one particular type of stimulus on one light cell so that the same input was being fed into the neural column at every cycle. This meant that the robot habituated to that input and therefore began to ignore the stimulus. This continued presentation of one stimulus until the robot stops responding to it is known as habituation training, and formed the first part of all the experiments performed. The second part involved the presentation of a second stimulus. This stimulus was presented

		Second Stimulus			
F		S1	S2	S3	S4
i	S1	0s	0s	0s	0s
r	S2	50s	0s	**3**s	0s
s	S3	51s	**2**s	0s	0s
t	S4	52s	37s	32s	0s

Table 1: Time taken for the robot to stop responding to a stimulus (rows) after complete habituation to a previous stimulus (columns). Note that the matrix is lower triangular except for the S3,S2 entry (highlighted). This shows that S2 and S3 occupy the same level of the hierarchy. Each result is the average of 10 trials (rounded to nearest integer value), although all measurements were within one second of the given result.

as soon as the robot ceased to move in response to the initial stimulus.

We performed two series of experiments. In the first we investigated the existence of the hierarchy of habituation. After habituating to one stimulus, a second was presented to the same light cell, and we recorded how the robot responded. The second set were aimed at investigating the locus specificity of the model, using experiments analogous to presenting prey dummies so that they appeared on different parts of the toad's retina. In our case we shone the light first onto one light cell and then onto another.

Each experiment was repeated a total of ten times and the mean average of the results (accurate to the nearest second) calculated.

6 Results

The results described in this section are a qualitative description of the behaviour of the robot. For a discussion of the quantitative results taken from a simulation of the network model, see section 7.

All the experiments took the same, two-part, form: one stimulus was presented a number of times and then a second stimulus was presented. Details of the experiments are given in the previous section. As mentioned in section 5, the time taken for the robot to stop responding to continuous presentations of a stimulus depends on which stimulus it is. It took 52 seconds to stop responding to stimulus S1, 37 seconds for S2, 32 seconds for S3 and only 20 seconds for S4.

6.1 Investigating the Existence of the Hierarchy

The first series of experiments concerned the behaviour of the robot when stimuli were presented to the same part of the retina each time.

Table 1 shows the number of seconds for which the robot responded to the stimuli presented as shown. The existence of the habituation hierarchy and its form can be inferred from the table. Training with S1 first means that none of the stimuli elicited any response, while presenting S4 first did not prevent the robot responding to the other stimuli. This tells us that S1 is above S4 in the hierarchy. The fact that S2 and S3 affect S4 puts them above it, but below S1. Presenting first S2 and then S3 is similar to S3 follow by S2. This puts them on the same level.

The habituation hierarchy is defined in the following way:
A pair of stimuli are at the same level in the hierarchy if their effects on each other are symmetrical, that is showing stimulus A followed by stimulus B is equivalent to showing stimulus B followed by stimulus A. Reasons for the existence of the hierarchy are discussed in section 8.

6.2 Locus Specificity

A number of tests were made to show that the observed behaviour was invariant over the light cells used and the stimuli presented. We observed that changing the light cell is equivalent to completely resetting the habituation mechanism, so if a stimulus is habituated on cell 1 and then another light cell is stimulated, however briefly, before returning the stimulus to the first cell again, it is as if the original habituation had never occurred.

7 Simulation

The FischerTechnik robots used here cannot download results. For this reason we also produced an isolated model of the network on which we could present values for arbitrary numbers of presentations. By using the values chosen for different lights on the robot and recording the output of the network we were able to produce the graphs shown in figures 5 and 6.

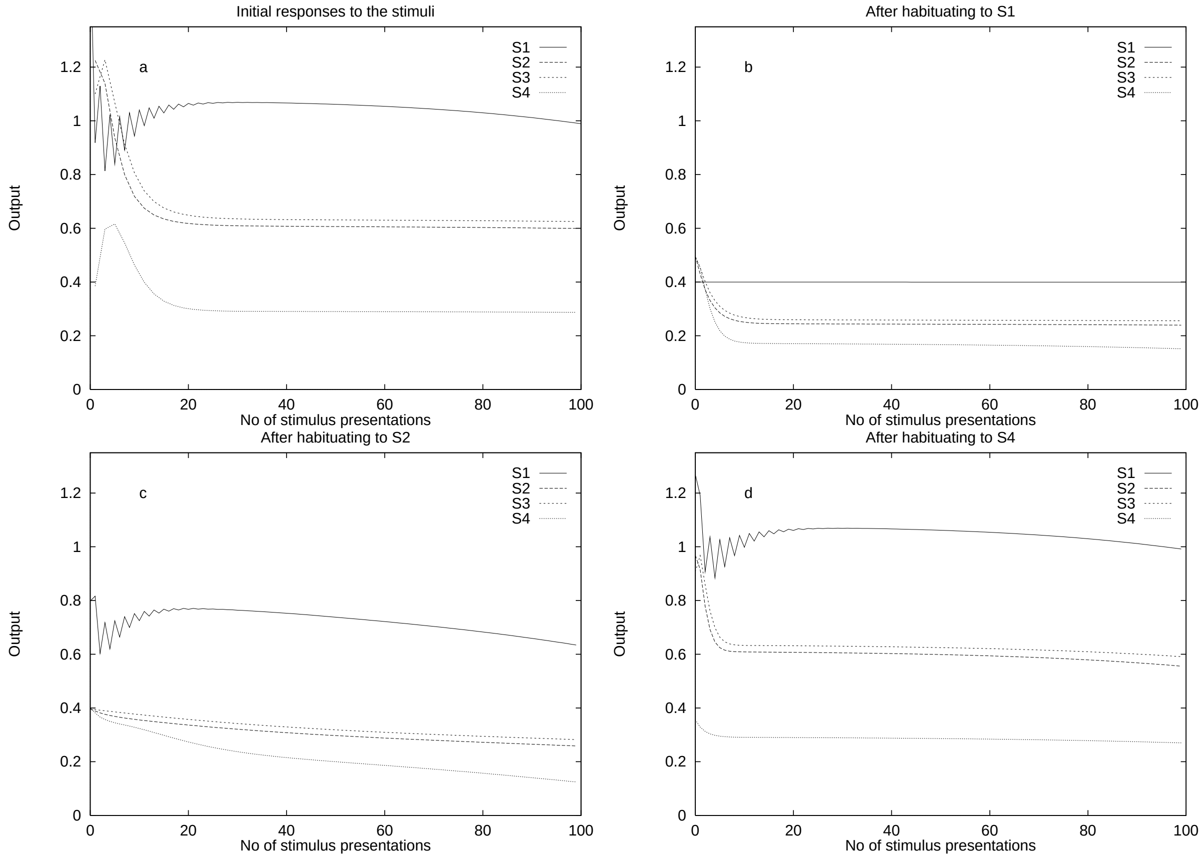

Figure 5: Graph ‘a’ shows the initial response of the network to presentations of each of the stimuli. The remaining graphs show the responses when each stimulus is presented after prior habituation to stimuli S1, S2 and S4 respectively. It can be seen that the responses from S2 and S3 are always similar. Habituating to S1 has a large effect on the other stimuli, whereas habituating to S4 has much less effect, with S2 being somewhere inbetween. This describes the hierarchy. The graph showing the effects of habituating first to S3 is not shown. It is the same as graph ‘c’ with the S2 and S3 lines reversed.

The graphs show the hierarchy clearly. By comparing graphs ‘a’ and ‘d’ we can see that habituating to S4 does not affect the response to the other stimuli at all, whereas habituating to S1 causes the response to be severely reduced. Due to frictional effects the robot stops turning when the output drops below about 0.4, since the motors do not receive enough power to drive them then. This is why no motion was seen on the robot.

8 Discussion

The last two sections have given qualitative and quantitative results obtained from the model with particular stimuli. It can be clearly seen that the model fulfils its primary objective of making the response of the robot to the presentation of a light stimulus habituate, so that the response decreases and eventually the robot stops moving towards the light. Other features such as the fact that stronger stimuli take longer to habituate are also demonstrated by the model. However, the most notable feature is the hierarchy of habituation. The hierarchy is observed in lower vertebrates, but not in mammals, which some think means that it is a relic of evolutionary development.

To investigate what caused the hierarchy we plotted the effects of using input signals $I(t)$ between 0.25 and 0.85, which are shown in figure 6. The outputs form three distinct bands of curves, with all the curves lying in one of the bands. This is the reason for the hierarchy is that whatever the input, the network model gives an output which falls into one of the three bands. The individual output lines are still visible within each band of curves, so within a level of the hierarchy some stimuli have slightly less effect than others, but these differences are very small.

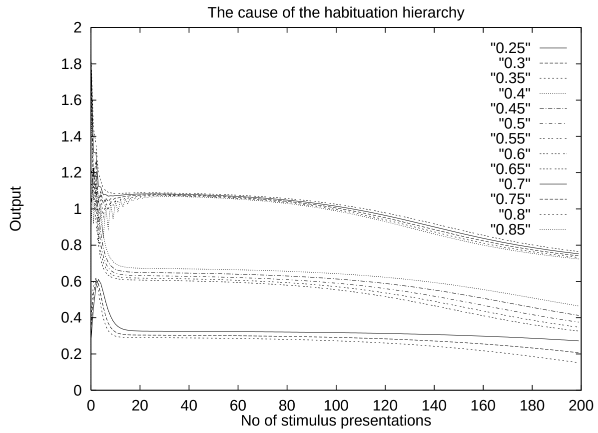

Figure 6: The graph shows the habituation provided by the network when input signals between 0.25 and 0.85 are used. The way that they divide naturally into three sections is the cause of the habituation hierarchy.

The very close similarity between S2 and S3 is demonstrated by the fact that they lie in the same band of the graph. In table 1 we saw that it takes slightly longer for S2 to habituate after training with S4 than S3 does. This indicates that, while the difference between them is small, it does still exist. However, if we had chosen S2 = 0.7 and S3 = 0.6, the hierarchy would have been very different, with two stimuli sharing the top level and one on each of the lower two levels. This is because the divide occurs between these two values.

While the experiments carried out in this paper are analogous to those mentioned in the biological literature, we are investigating only the effect of the model of the medial pallium. This is only one part of the brain of the toad, there are several sections preprocessing the signals received from the eyes, and responding to the habituated signal. As far as possible we have simplified all other sections to concentrate on the habituation model, so we are using the four very basic light sensors which record only the intensity of the light as an extremely primitive 'eye'. This is obviously much simpler than the eyes of a toad.

9 Conclusions and Future Work

This paper shows that a suggested model which displays a habituation mechanism can be transferred onto a mobile robot so that the robot eventually stops responding to repeated stimuli. This model includes a hierarchy of habituation. In the implementation presented here, which was designed purely to test the model, sensory perceptions were known in advance and split into just four categories. This would certainly not be the case for a robot moving in the real world, so an investigation of how inputs should be prepreprocessed is needed.

Animals use habituation for several purposes. The primary use is to enable the animal to ignore weak and repeated stimuli so that it can concentrate on other, potentially more interesting or dangerous stimuli. This effect has applications in the field of robotics, a robot moving in a 'noisy' environment, where there are lots of extraneous stimuli (such as flashing lights) which distract the robot from performing its task, could use the habituation mechanism to learn to ignore those stimuli which did not affect its purpose.

However, the initial motivation for the study lay in another direction. Applications of a robot which could detect new or novel things within an environment are myriad. Unfortunately, the question of what is novel is an exceedingly hard one. A different, and more tractable approach to the problem is to define normality as 'things which you have seen before' and to learn what is normal for a particular environment. One possible way of doing this is to ignore things which have been observed before, so that only new features are noticed. This is exactly what the habituation mechanism does. The challenges for this approach lie in learning about how the features should be presented to the habituation module.

The model investigated here is of interest because of its biological plausibility. However, for use in a robot performing other tasks it has several disadvantages, principally the computational cost of using it.

The hierarchy of habituation which has been demonstrated in the experiments discussed previously is an interesting side–effect of the model, but one whose applications to robotics are not obvious to us. In some cases it may be useful to have a habituation mechanism that, on habituating to one stimulus, automatically ignores stimuli considered less important, although how to decide which stimuli are important is an open question. We could not find references in the biological literature of a real use of the hierarchy.

Future work will include an analysis of the model to find ways in which it can be simplified while retaining the desired behaviour.

References

[1] Bailey CH, Chen MC (1983) Morphological Basis of Long-Term Habituation and Sensitization in *Aplysia*. Science 220:91-93

[2] Byrne J, Gingrich K (1989) Mathematical Model of Cellular and Molecular Processes Contributing to Associative and Nonassociative Learning in Aplysia. In Byrne J, Berry W (eds) Neural Models of Plasticity. Academic Press, pp 58-72

[3] Castellucci VF, Carew TJ, Kandel ER (1978) Cellular Analysis of Long-Term Habituation of the Gill-Withdrawal Reflex of *Aplysia California*. Science 202:1306-1308

[4] Ewert J-P, Kehl W (1978) Configurational Prey-Selection by Individual Experience in the Toad Bufo Bufo. Journal of Comparative Physiology A 126:105-114

[5] Greenberg S, Castellucci V, Bayley H, Schwartz J (1987) A Molecular Mechanism for Long-Term Sensitization in *Aplysia*. Nature 329:62-65

[6] Groves PM, Thompson RF (1970) Habituation: A Dual-Process Theory. Psychological Review 77: 419-450

[7] Hawkins R, Kandel D, Siegelbaum S (1993) Learning to Modulate Transmitter Release: Themes and Variations in Synaptic Plasticity. Annual Review of Neuroscience 16:625-665

[8] Kohonen T (1989) Self-Organisation and Associative Memory, 3rd ed. New York: Springer-Verlag

[9] Lara R, Arbib MA (1985) A Model of the Neural Mechanisms Responsible for Pattern Recognition and Stimulus Specific Habituation in Toads. Biological Cybernetics 51:223-237

[10] Stanley JC (1976) Computer Simulation of a Model of Habituation. Nature 261:146-148

[11] Thompson RF (1986) The Neurobiology of Learning and Memory. Science 233:941-947

[12] Thompson RF, Spencer WA (1966) Habituation: A Model Phenomenon for the Study of Neuronal Substrates of Behaviour. Psychological Review 73:16-43

[13] Wang DL, Arbib MA (1991) How Does the Toad's Visual System Discriminate Different Worm-Like Stimuli? Biological Cybernetics 64:251-261

[14] Wang DL, Arbib MA (1991) Hierarchical Dishabituation of Visual Discrimination in Toads. In Meyer J-A, Wilson S (eds) Simulation of Adaptive Behaviour: From Animals to Animats. MIT Press, Cambridge, pp 77-88

[15] Wang DL, Arbib MA (1992) Modelling the Dishabituation Hierarchy: The Role of the Primordial Hippocampus. Biological Cybernetics 76:535-544

[16] Wang DL, Ewert J-P (1992) Configurational Pattern Discrimination Responsible for Dishabituation in Common Toads *Bufo bufo* (L.): Behavioural Tests of the Predictions of a Neural Model. Journal of Comparative Physiology A 170:317-325

[17] Wang DL, Hsu CC (1990) SLONN: A Simulation Language for the Modelling of Neural Networks. Simulation 55:69-83.

[18] Wang DL (1993) A Neural Model of Synaptic Plasticity Underlying Short-term and Long-Term Habituation. Adaptive Behaviour 2:111-129

[19] Wang DL (1995) Habituation. In Arbib MA (ed) A Handbook of Brain Theory and Neural Networks. MIT Press, Cambridge, pp 441-444