A Robot Implementation of a Biologically Inspired Method for Novelty Detection

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

5.4.2001

Abstract

This work examines the ability of a biologically inspired novelty detection model to learn and detect changes in the environment of a mobile robot. The novelty detection model used was inspired by recent neurological findings of novelty neurons in monkeys' perirhinal cortices. Experiments examine the difference required between stimuli before the novelty detection model recognises them as novel and the ability of the model to learn its environment on-line. The novelty detection model examined in this paper is based on calculating the energy of a Hopfield network. It appears to be potentially useful for on-line learning on mobile robots as it can reliably learn from a single presentation of a novel stimuli. A qualitative comparison is made to an alternative model that also carries out novelty detection on a mobile robot.

1 Introduction

The ability to detect and respond to changes in the environment would intuitively appear to be advantageous to any agent. Indeed studies from Pavlov (1927) onwards have identified an involuntary mechanism that is capable of drawing awareness to significant changes in the environment even if voluntary attention is currently focused elsewhere. This skill would be useful to a mobile robot; minimising the computational effort required to evaluate all the incoming stimuli from its environment, and focusing its attention on items that are 'significant' - either not seen before, not seen recently or particularly salient to the agent. From this it can be seen that novelty detection or familiarity discrimination could provide a part of a 'involuntary attention' mechanism.

The medial temporal lobe and especially the perirhinal cortex have been implicated by many studies (review by Brown and Xiang (1998)) as being necessary for familiarity discrimination of visual stimuli in monkeys. A study of electro-physiological recordings from neurons in this region (Xiang and Brown, 1998) identified neurons whose response appears to encode detection of novelty among the presented stimuli. This study was carried out on two monkeys (*Macaca mulatta*). These *novelty neurons* were found throughout the anterior inferior temporal cortex, including perirhinal cortex, and also entorhinal cortex. They were identified by their first and repeat reactions to familiar[1] and novel[2] stimuli during a recording session. They respond strongly to the first presentation of a novel stimulus then only weakly to repeat presentation of the same novel stimuli some 4-8 minutes later. The novelty neurons showed even less response to familiar stimuli. It is suggested (Xiang and Brown, 1998) that novelty neurons form part of the mechanism that allows primates to perform familiarity discrimination.

A simulated spiking neural network model that closely replicates the recorded output response of novelty neurons has been developed by Bogacz et al. (1999b, 2000). The computation performed by this

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

[1]shown to the animals each day

[2]used twice a day and not again for at least 2 months

model in determining familiarity can be shown to be very similar to that performed in calculating the energy of a Hopfield network (Bogacz et al., 1999b, 2000). Evaluating the energy of a Hopfield network is a simple algorithm whose execution time is fixed (irrespective of the number of patterns stored), making it an ideal model for on-line operation in a mobile robot.

2 Model Details

The model stores information about familiar patterns (patterns that it has learnt) in the weights of a Hopfield network. Classification of new patterns is determined by calculating the energy of the network for the pattern to be classified. Patterns with low energy are generally familiar, those with higher energies are generally novel. The model sacrifices the ability of the Hopfield network to retrieve or complete previously learnt patterns, but the benefit of this tradeoff is that it can classify significantly more patterns as familiar or novel than a Hopfield network can typically recall.

In the simplest model the Hopfield auto associative network used is a fully connected recurrent network of N neurons. Each neuron can take the value $\{+1, -1\}$; $+1$ representing an active state, -1 an inactive state. Each pattern $\boldsymbol{\xi}^\mu$ stored in the network is also N bits long, where the i^{th} bit of the pattern is given by ξ_i^μ. In order to store P patterns in the network, the network weights w_{ij}, are computed as:

$$w_{ij} = \begin{cases} \frac{1}{N}\sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu & \text{if } i \neq j \\ 0 & \text{if } i = j \end{cases} \tag{1}$$

For on-line learning where patterns are learnt one at a time this becomes: $w'_{ij} \leftarrow w_{ij} + \frac{1}{N}\xi_i\xi_j$ if $i \neq j$. If an arbitrary pattern $\boldsymbol{x}$ is presented to the network, where x_i and x_j are the i^{th} and j^{th} bits of this pattern, then the energy of the network is defined (Hopfield, 1982; Hertz et al., 1991) as:

$$E = -\frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} w_{ij}x_ix_j \tag{2}$$

Normally the Hopfield auto associative network is used as a content addressable memory. In this mode of operation a partial pattern which is to be recalled is presented to the network and then the state of each neuron in the network is updated several times until the network *relaxes* to the recalled pattern. Recall tends to take several cycles through all the neurons in the network until none of the neurons change state, or some arbitrary number of cycles have been completed. By contrast calculating the energy of a Hopfield network only requires the calculation of the above summation – equation 2. Thus execution time of this algorithm is fixed.

Familiarity discrimination is achieved by checking the energy of a pattern against some threshold. It is possible to show (Bogacz et al., 1999a, 2000; Crook, 2000) that the energy E for a pattern which has been learnt by the Hopfield network is:

$$E \quad \approx \quad -\frac{N}{2} + \theta\left(0, \sqrt{\frac{P}{2}}\right) \tag{3}$$

and the energy for a *novel* random pattern is:

$$E \quad \approx \quad \theta\left(0, \sqrt{\frac{P}{2}}\right) \tag{4}$$

Where N is the number of neurons in the network, θ is a noise term modelled by a Gaussian distribution with mean of zero and standard deviation $\sqrt{P/2}$, and P is the number of patterns stored in the network. Based upon this a threshold of $E < -N/4$ is used for classification of patterns.

By recognising the symmetrical nature of the Hopfield network's weights, $w_{ij} = w_{ji}$, and that $w_{ij} = 0$ when $i = j$ the total number of terms that need be computed can be reduced from N^2 to $\frac{1}{2}(N-1)N$, saving both memory and processing time. As only half the terms of equation 2 are computed the threshold for familiarity classification is halved, i.e. $E < -N/8$.

With reference to equations 3 and 4, for the Hopfield model to accurately classify novel or familiar patterns the noise term θ needs to be small. If the probability of the model making a *recognition* error is constrained it can be shown that the maximum number of patterns that can be stored increases in proportion to N^2 (Bogacz et al., 1999a, 2000; Crook, 2000). By contrast the maximum number of patterns that can be stored and accurately *recalled* by a Hopfield network is proportional to N (Hertz et al., 1991).

3 Robotic Implementation

Analysis of the Hopfield model's properties (Crook, 2000) suggest that in order for it to demonstrate a performance advantage over simpler techniques, such as just storing every learnt pattern, the model needs to be used with a sensor array that has a large numbers of bits per pattern, and the model needs to deal with very large numbers of patterns. In order to achieve input patterns that contain a large number of bits and also the possibility of generating large numbers of patterns, it was decided to use video images captured from a camera. The camera used is mounted on a B21 mobile robot, see figure 1. Using a robot mounted camera allows the model to carry out online learning while the robot manoeuvres through an environment.

The colour video image from the camera is captured and then processed. The processing consists of extracting from the image those pixels corresponding to a particular colour, in this case orange. This results in a much simplified image where only orange objects are seen and the rest of the world is ignored. Orange was selected as it appeared, with the colour video camera used, to be the most invariant to changes in the lighting of the laboratory. Simplifying the image in this way gave fairly straight forward control over what the robot perceived. Issues relating to more biologically realistic pre-processing of images; dealing with scale invariance, or known visual processes such as detecting line features, was consider to be outside the scope of this work – the video images were primarily considered to provide a large binary array of sensory input with which to test the network's function.

The size of the video image determines the number of neurons in the Hopfield network and vice versa. Consideration of the available memory on the robot gave the following parameters for the model: post processed binary image of 48×48 bits, Hopfield network of $2,304$ neurons, storage of $2,653,056$ weights (requiring 10.6 Megabytes of memory), threshold for novelty of $-\frac{N}{8} = -288$, theoretical number of patterns that can be stored if the error rate is not to exceed 1% is $P_{max} = 0.023N^2 = 122,093$. It was found that classification of a pattern by this model took around 1.1 seconds – running on a 100MHz Pentium with 32 Megabytes RAM. Updating the weights to learn a pattern took a similar time.

4 Experiments

The objective of the experiments were to explore the effectiveness of the Hopfield model as a method of determining novelty and learning in a mobile agent. The design of the second set of experiments, section 4.2, was such that it is possible to make some qualitative comparisons with the experiments carried out by Marsland et al. (2000) using a Habituating Self-Organising Map (HSOM) network.

4.1 Threshold Test

Experiments were conducted to confirm the theoretical observation (Crook, 2000) that patterns need to differ by at least 15% to be recognised as novel. The robot is shown a rectangular piece of orange card. It learns this image and then proceeds to back away from the card until the change in the image is such that it appears novel. The total number of pixels that differ between the initial and final images is then calculated. This was repeated three times to confirm that the result was consistent.

Some limited tests were also carried out to explore how learning other patterns effects the ability of the model to detect novelty. The robot is trained on an increasing number of other images and the effect of this on the threshold test above is observed. As each new pattern is introduced it is first classified by the model to confirm that it appears novel when compared to the set of patterns already learnt. The orange rectangle is then shown to the robot which backs away while still viewing the card.

4.2 On Line Learning

In Marsland et al. (2000) a robot travelled the length of a corridor while a HSOM network learnt and classified the patterns detected by the robot's sonar sensors. In a parallel approach to this, the B21 robot travelled along one wall of a laboratory while classifying and learning patterns detected by the video camera, see figure 1. The robot is started with an untrained Hopfield network and alternating 'learning' and 'non-learning' runs are made along a 'gallery' of 'pictures'. Once the robot no longer classifies any element of the gallery as novel, the pictures are modified and the robot retraces its route along the length of the gallery while classifying the new perceptions. The images presented to the network are not invariant to changes in distance from the camera. To compensate for this the starting position of the robot and its distance from the wall as it travels along the galley is maintained by a simple wall following algorithm. As the robot travelled along the gallery plots were recorded showing the change in the energy level of the Hopfield network and the points where the model classified the image it saw as novel. The aims of these tests were to demonstrate if the novelty detection model could: (i) initially learn its environment (ii) recognise changes in the environment (ii) learn these changes, (iii) no longer regard any changes as novel once it has been allowed to learn them.

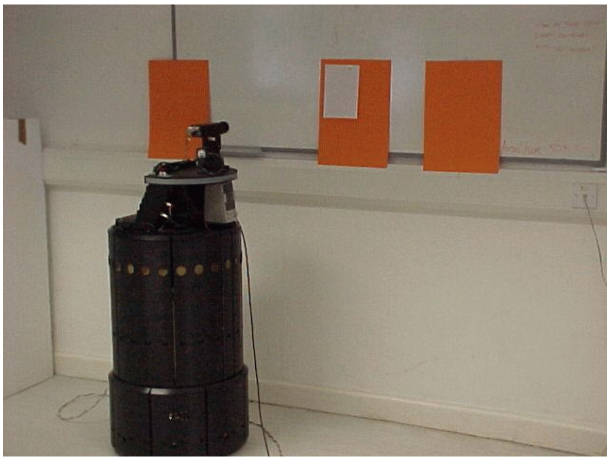

Figure 1: B21 Robot and Gallery

5 Results

5.1 Threshold Test

As can be seen from the table 1 and figure 2, when only the rectangle has been learnt, the percentage change in the image when it is no longer classified as familiar it is close to the 15% predicted[3]. This result can be seen as stable for all three runs. The learning of other patterns increases the percentage change required before the rectangle is recognised as novel. The change required rising to 25% with the addition of only two patterns. The third run failed because the robot appeared to recognise the shrinking image of the orange rectangle. After comparing the shape it could see with the patterns learn it turned out that the shrinking rectangle differed by only 14.7% from pattern three that had just been introduced. The Energy curves for all these runs are shown in figure 2.

5.2 On Line Learning

Figures 3 to 7 show the various states of the gallery in chronological order. Under each gallery the various runs that were made are shown. For most states of the gallery three runs are shown, one when the robot looked at what had changed (without learning), a second run where it was allowed to learn and a third

[3]The percentage change was measured by comparing the two post-processed images as presented to the network and counting the number of pixels that they differed by.

Run	Pattern Added	Energy of Added Pattern (before learning it)	Start Distance (mm)	Starting Energy of Test Pattern	End Distance (mm)	End Energy of Test Pattern	Change in Image
1	—	—	651	-573.8	741	-266.5	16.1%
2	—	—	653	-573.8	730	-281.0	14.9%
3	—	—	649	-572.8	730	-273.4	15.6%
4	1	-284.5	654	-856.3	766	-284.7	18.9%
5	2	-153.5	655	-932.6	810	-287.3	24.7%
6	3	-100.3	656	-948.8	1104	-308.9	—

Table 1: Percentage change in test pattern before it is no longer recognised (a pattern is seen as novel if its energy is above the threshold of -288).

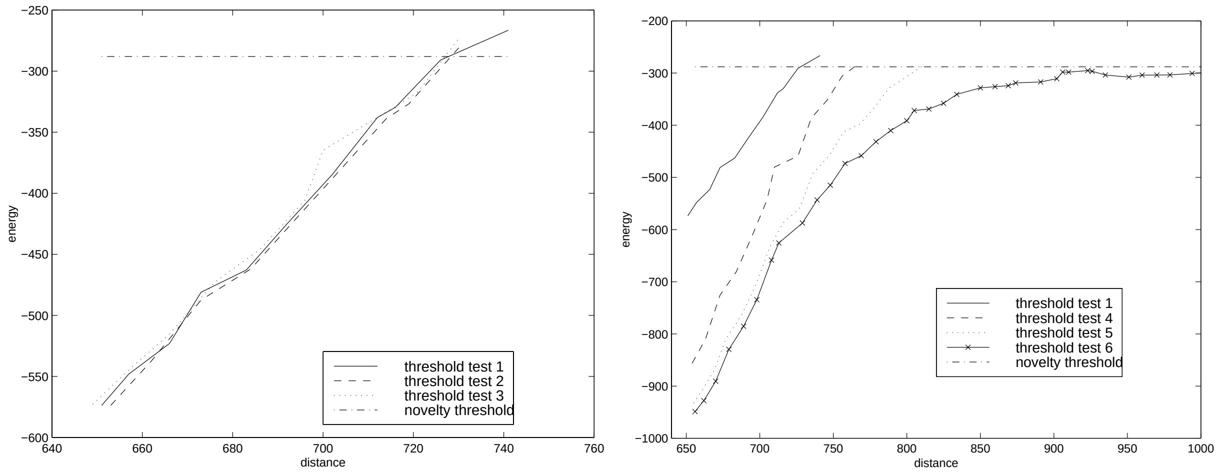

Figure 2: Energy Curves for Threshold Tests; left to right (i) only test pattern learnt, (ii) additional patterns incrementally learnt

run which demonstrates if the learning is effective. What the robot can see at any one time can be visualised by mentally tracking the dotted window shown in each figure along the gallery.

Figure 3 shows the first two runs made by the robot. Initially the network was untrained, all its weights are zero. This gives an energy of zero for all images and given that the threshold for novelty is -288 everything the robot looks at will be regarded as novel. The robot's first action is therefore to learn the first image that it is presented with, this is indicated by the immediate drop in the energy curve from 0 to -575.75 at the start of the first run plot. As the robot tracks along the wall the image that it can see becomes gradually less and less like the first image that it learnt. This is indicated by the rise in the energy level until it reaches the novelty threshold. At this point the network learns this new image, adding it to the one it has already learnt. These actions are repeated at different positions until the robot has learnt four views of the orange card that it first saw. It learns again just as this card is at the far left hand side of the robot's field of vision. This corresponds to the point where the robot observes for the first time the largely featureless section between the first and second cards. The energy is then fairly static until the next card begins to fill the robot's field of vision. The energy level initially rises to just over -500 before falling sharply as the second orange card tracks towards the left of the robots field of vision. It would appears that this corresponds to the position in the robot's field of vision that the first card occupied when its image was learnt. The energy level then rises again until the narrow gap between the second and third card occupies the centre of the robot's vision. At this point it decides this is novel and learns it. It learns one final time as this same gap approaches the far left hand side of the image.

The second run was made with learning disabled to see how much of the gallery remained novel. As

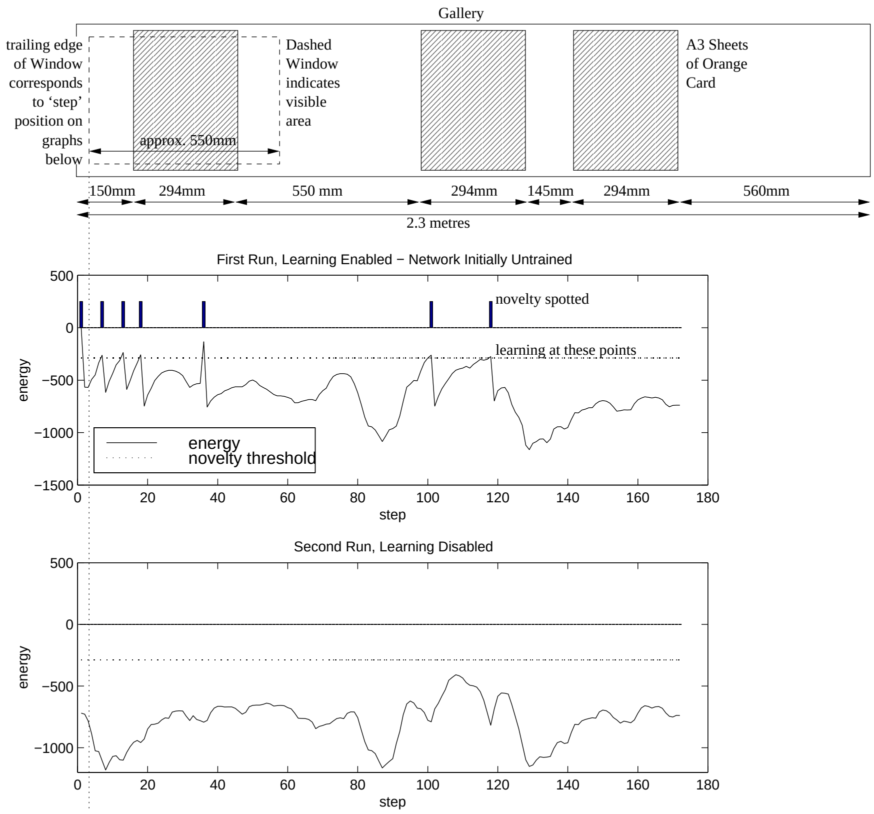

Figure 3: Initial Gallery

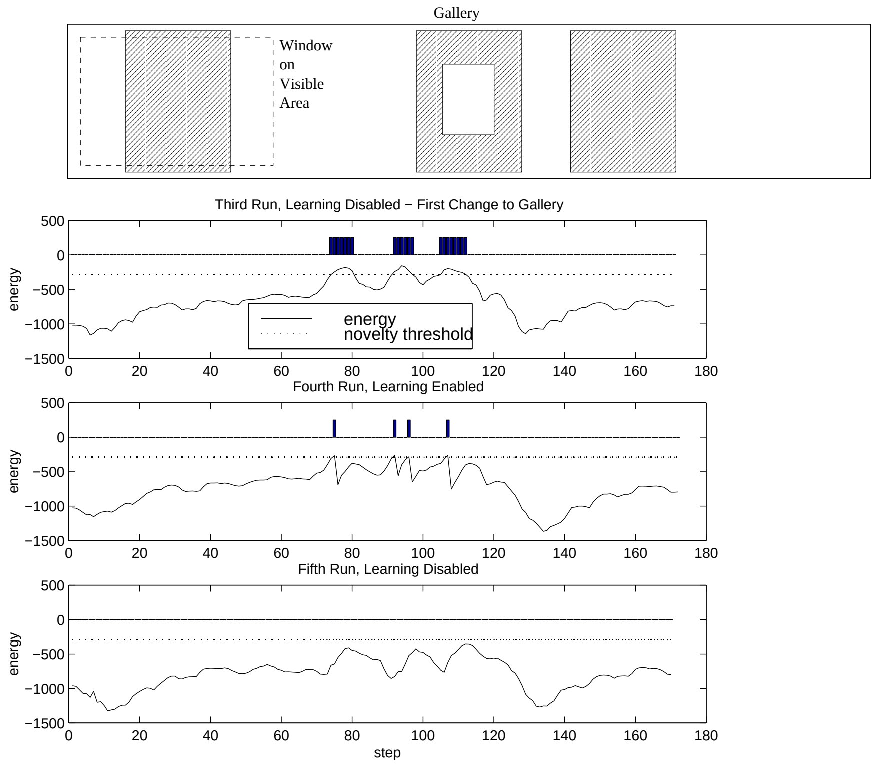

Figure 4: First Change to Gallery

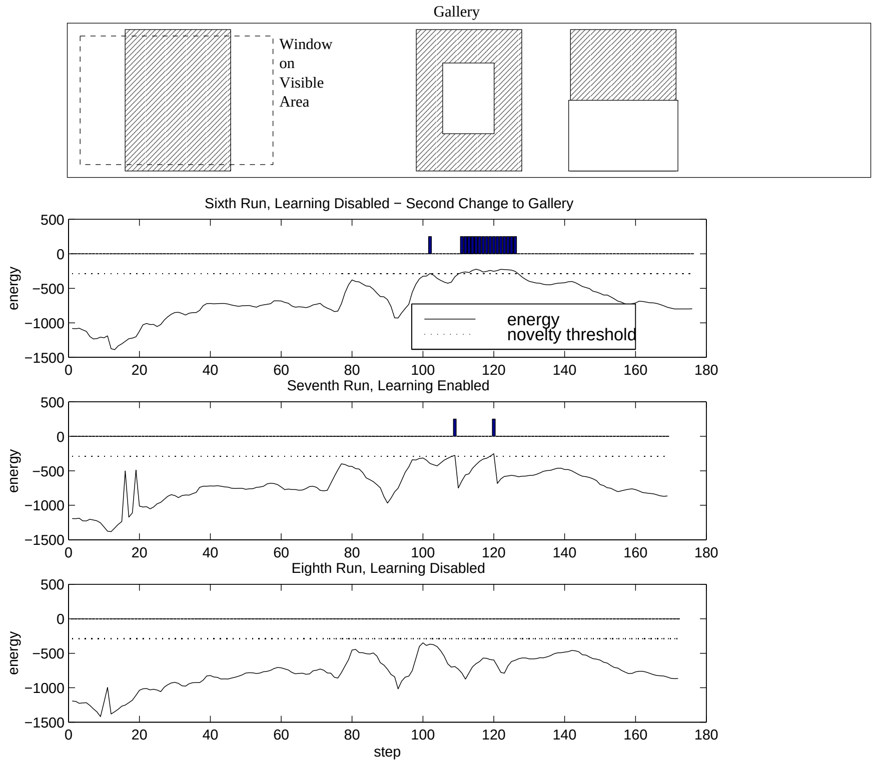

Figure 5: Second Change to Gallery

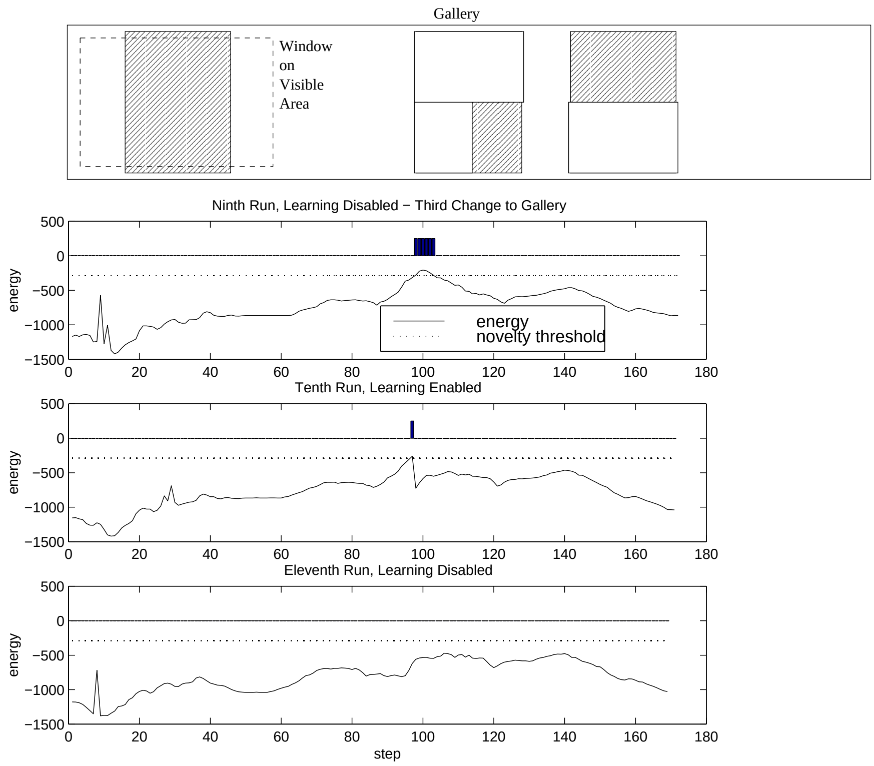

Figure 6: Third Change to Gallery

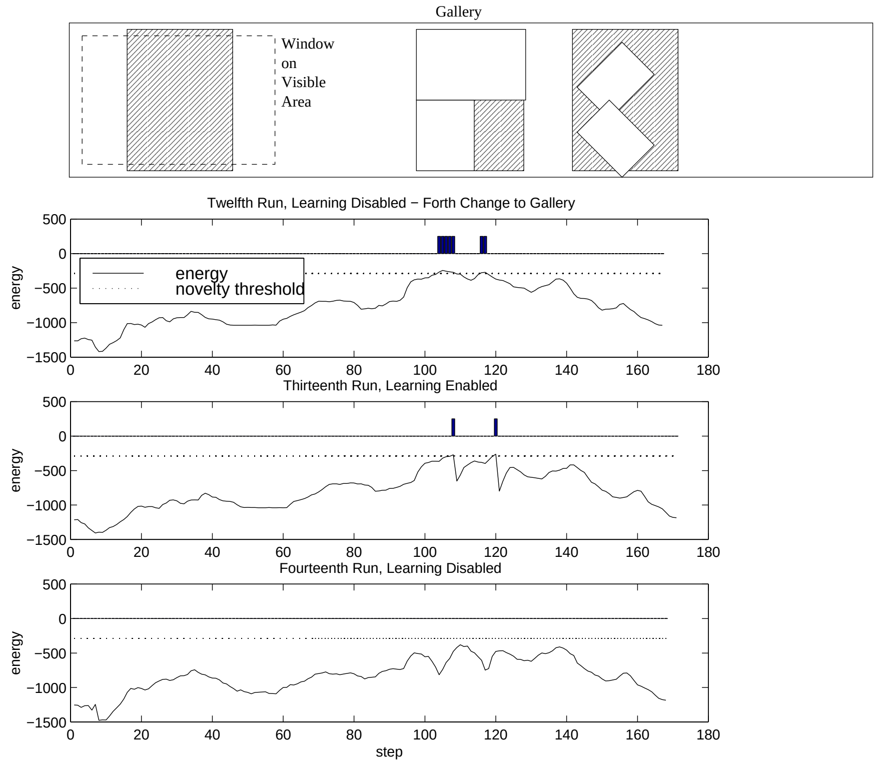

Figure 7: Fourth Change to Gallery

can be seen in figure 3 all of the features were now regarded by the robot as familiar. Looking at the energy curve it can be observed that where learning occurred during the first run the energy tends to dip. This is especially true for each of the rectangles of card, where the effect of learning the image of the first card some four times results in a significant drop in the value of the energy. The narrow gap between the second and third cards remains of interest as indicated by the peak in the energy curve, although the maximum of the peak is insufficient for the gap to be still classified as novel. Either side of this peak are dips — results of the learning that occurred when the robot tracked past the gap in the first run.

The gallery was then changed by placing a sheet of A5 in the centre of the second card, as shown in figure 4. During the third run learning was disabled so the robot tracked along the gallery just reporting on what it saw. As can be seen it classified the change to the second card as novel three times; (i) as soon as it came fully into view, (ii) again when it was slightly to the left hand side of the robot's field of view — the position where the energy normally dips for the unadulterated rectangles of card and (iii) finally when both the change to the second card and the gap between the second and third card were in view. Dips in the energy curve can still be observed when the first and third cards are seen by the robot.

Learning was enabled for the fourth run, and the robot learnt four times at around the same positions as described above. In the fifth run the robot has learn the new arrangement and no longer finds it novel. The energy curve still peaks around the position of the modified second card but the maxima are not significant enough for the images to be classified as novel.

In the sixth run the gallery was modified by placing a sheet of A4 over the bottom of the third card, see figure 5. The robot finds this change novel as soon as it appears on the right hand side of its field of view and the energy remains high as the doctored card tracks to the centre of the field of view. When learning is enabled in run seven the robot does not react to this change quite as early as in run six. This is probably due to variation in the distance between the robot and the gallery. It learns twice: (i) first when it can see around three quarters of the change to the third card, part of the previous change to second card and the gap between the two cards. (ii) The second time is when the change is central in the field of view. The eighth run shows it has again successfully learnt this new arrangement. Further changes were made to the gallery, figures 6 and 7. Each time the robot successfully demonstrated that it could detect and then learn these new arrangements.

Some spikes in the energy appear near the start of runs seven, eight, nine, ten and eleven. They appear at different points each time and there is no obvious explanation as to their cause.

6 Discussion

When no other images have been learnt the percentage change required for an image to be classified as novel is around the theoretical prediction of 15%. Interference caused by learning other images does appear to increase increases the percentage change required although further work is required to quantify this effect.

The model learns an image after a single exposure, and reliably recognises it on subsequent runs despite possible noise in the video image and inaccuracy in positioning of the robot. The model is able to recognise all elements of the gallery after learning only a few patterns. Changes, once they were identified, can be reliably learnt and will then be perceived as familiar. Changes to the gallery can be detected provided they are significantly different – there is suggestive evidence that the percentage difference might need to be greater than the identified minimum of 15%. This may be due to the properties of the Hopfield network model or an artifact of the images and pre-processing employed, i.e. the input stimuli could be cluster in a small part of the possible stimuli-space.

In examining the energy curves from the various runs along the gallery, it is interesting to note that some patterns seem to retain some higher level of novelty than others. For example the gap between the second and third cards has a relatively high energy even after being learnt from several positions. The presences of such high points in the energy curve suggests that they would help promote the energy level of nearby changes, i.e. placing a change alongside the gap makes it easier for the novelty filter to decide that it is novel. Therefore it is not only the magnitude of the change that is important but also its relative position in the environment.

Comparisons can be made to the tests carried out by Marsland et al. (2000). In those tests a robot travelled a distance of 10 meters down sections of corridor, while every 10 cm presenting sonar perceptions to a novelty filter. As in section 4.2 this robot made alternating learning and non-learning runs and retained the network weights learnt during previous runs. The results appear very similar. In both cases the robots were able to successfully learn their environment and then perceive novelty when the environment is changed. Both also were able to learn changes in the environment so that they no longer perceived them as novel. There are however a few difference in results. In the results presented by Marsland et al. (2000) when the robot was completely untrained it took three runs with learning enabled before it ceased to detect novelty. Similarly when the environment was changed Marsland et al. (2000)'s robot took two runs with learning enabled before it ceased to detect the change. This compares to the one run required by the Hopfield model, to both initially learn its environment and then update its knowledge when a change occurred. More speculatively it appears that the Habituating Self-Organising Map (HSOM) network used by Marsland et al. (2000) may pick out finer details than can be detected by the Hopfield model. Marsland et al. (2000) report that their robot periodically detected a crack in the wall of the corridor which is very thin. However no details are given on how large an impact this crack has on the sonar stimuli received by the network. The relative sensitivity of the two methods may however depend on differences in the source of the stimuli used. It is possible that the stimuli produced from sonar data is more varied than the data produced by the vision system, i.e. the patterns are less clustered within their respective stimuli-space.

7 Conclusions & Further Work

Overall the Hopfield base model appears a potentially useful model for a novelty detector, especially as it appears to be able to reliably learn from a single presentation of a novel pattern. The experiments demonstrate that the novelty detection model recognises changes in the gallery and can learn these changes. The model can learn the entire gallery in a single run, compared to several for the HSOM network used by Marsland et al. (2000).

Further work is required to: (i) Quantify the effect that learning patterns has on the ability of the Hopfield network model to distinguish between novel and familiar patterns. (ii) Establish *quantitatively* the relative sensitivity of this and other novelty detection model and determine if any apparent insensitivity in the Hopfield model is a property of the model or of the selection of images and pre-processing used.

References

Bogacz, R., Brown, M. W., and Giraud-Carrier, C. (1999a). High capacity neural networks for familiarity discrimination. In *Proceedings of ICANN, Edinburgh*, pages pp773–778.

Bogacz, R., Brown, M. W., and Giraud-Carrier, C. (1999b). Model of familiarity discrimination in the brain - efficiency, speed and robustness. In *EmerNet Workshop*, pages 23–26.

Bogacz, R., Brown, M. W., and Giraud-Carrier, C. (2000). Model of familiarity discrimination in the perirhinal cortex. *to be published in Journal of Computational Neuroscience*.

Brown, M. and Xiang, J. (1998). Recognition memory: Neuronal substrates of the judgement of prior occurrence. *Progress in Neurobiology*, 55:149–189.

Crook, P. (2000). Spotting novelty: A neural network model for familiarity discrimination. Master's thesis, Division of Informatics, Unvirsity of Edinburgh.

Hertz, J., Krogh, A., and Palmer, R. G. (1991). *Introduction to the theory of Neural Computation*. Perseus Books.

Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. In *Proceedings of National Academy of Sciences*, volume 79, pages 2554–2558.

Marsland, S., Nehmzow, U., and Shapiro, J. (2000). Detecting novel features of an environment using habituation. In *Proceedings Simulation of Adaptive Behaviour*. MIT Press.

Marsland, S., Nehmzow, U., and Shapiro, J. (2000). Novelty detection on a mobile robot using habituation. In *From Animals to Animats: The 6th International Conference on Simulation of Adaptive Behaviour (SAB'2000)*, pages 189 – 198. MIT Press.

Pavlov, I. P. (1927). *Conditioned reflexes*. Clarendon Press, Oxford.

Xiang, J.-Z. and Brown, M. (1998). Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology*, 37:657–676.