Vision-based Environmental Novelty Detection on a Mobile Robot

Stephen Marsland, Ulrich Nehmzow and Jonathan Shapiro
Department of Computer Science
University of Manchester
Oxford Road
Manchester M13 9PL, U.K.
{smarsland, ulrich, jls}@cs.man.ac.uk

Abstract

Novelty detection is about recognising features that do not fit into the pattern of previous perceptions. It is an important survival trait for animals, and is also useful for robots. For example, a robot equipped with the ability to detect novelty can select which features of an environment to investigate and learn about.

In this paper we enable a mobile robot to learn a model of an environment that the robot experiences through the images of a monochrome camera while exploring. Once the robot has learnt a model of this environment we move the robot to a new environment and ask it to detect those features of the new environment that do not fit into the model, i.e., the novel features. We describe a number of different algorithms for producing an input vector from the image that is suitable for presentation to the novelty filter, and demonstrate results using the approach that worked best.

1 Introduction

In animals, the ability to detect novel features of their environment is an important survival trait, as any unusual perception may be a potential predator. For robots, too, this ability can be very useful. For example, a robot that could detect novelty would be able to decide which features of an environment to concentrate on and learn about. Another use for a novelty detecting robot is as an inspection agent. In a training phase, the robot could explore a number of environments that were known to exhibit no problems. Once the robot had learnt an accurate model of these environments, it could explore the whole environment, highlighting those features that did not fit into the model previously acquired. These would be the sites of potential problems.

In previous work we have proposed a novelty filter that can operate online, and is therefore suitable for use on a mobile robot. The filter has been used to direct the attention of a mobile robot to novel and therefore potentially interesting stimuli [6], and to learn a model of an environment using the robot's sonar sensors [7]. In this paper, the novelty filter is used in conjunction with a monochrome camera to demonstrate that the concept can be extended to using camera images of an environment. The camera is positioned facing the wall nearest to the robot, and records an image every 10 cm of travel.

2 Related Work

A number of novelty detection techniques have been proposed in the literature. The first was Kohonen's Novelty Filter [5, 4], which is trained to reproduce the inputs at the output so, after training, any input presented to the network produces one of the learned outputs. Then, the bitwise difference between the input and the output highlights novel components of the input.

Several other approaches have used the Self-Organising Map [4]. Ypma and Duin [13] proposed a novelty detection mechanism based on measuring the quality of the match between a trained SOM and a particular dataset. By training the SOM on data that is known to be normal and then evaluating the measures on a new dataset, they compute how likely it is that the new dataset comes from the same distribution as the training data.

Another approach is to calculate the distance of the winning neuron from neighbourhoods that fired when training data known to be normal was introduced, and counting as novel those inputs where the distance is beyond a certain threshold. This method was used by Taylor and MacIntyre [12] to detect faults when monitoring machines. The network was trained on data taken from machines operating normally and data deviating from this pattern was taken as signifying a machine fault.

Another novelty filter, the FamE model, has been proposed by Bogacz *et al.* [1]. Inputs are presented to a Hopfield network, and the energy of the network evaluated. If the energy is low, the input is

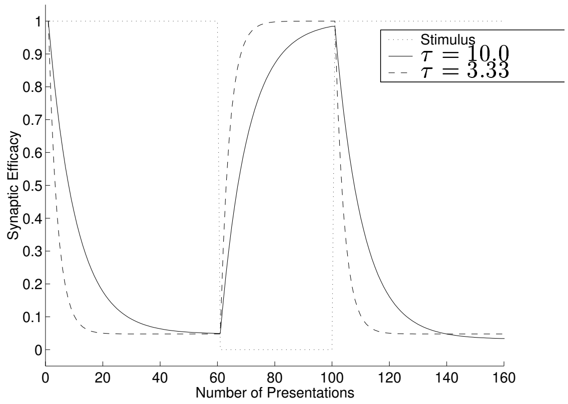

Figure 1: An example of how the synaptic efficacy varies when habituation is modelled using equation 1. In both curves, a constant stimulus $S(t) = 1$ is presented at $t = 0$, causing the efficacy to fall. The stimulus is removed ($S(t) = 0$) at time $t = 60$ where the synaptic efficacy rises again, and becomes $S(t) = 1$ again at $t = 100$, causing another drop. The two curves show the effects of varying τ in equation 1. A larger value of τ causes both the learning and forgetting to occur faster. The other variables were the same for both curves, $\alpha = 1.05$ and $y_0 = 1.0$.

assumed to be familiar, otherwise it is considered to be novel. The model has recently been applied to a mobile robot application [2]. The robot takes pictures of a 'picture gallery' of orange rectangles on a white wall and the FamE model evaluates the novelty of the simple images produced.

3 The Novelty Filter

3.1 Habituation

Habituation is a decrement in response when a stimulus is seen repeatedly without any ill effects. A number of researchers have produced mathematical models of the effects of habituation on the efficacy of a synapse. We use a simple model proposed by Stanley [11]. In his model the synaptic efficacy, $y(t)$, decreases according to the following equation:

$$\tau \frac{dy(t)}{dt} = \alpha \left[y_0 - y(t) \right] - S(t), \qquad (1)$$

where y_0 is the original value of y, τ and α are time constants governing the rate of habituation and recovery respectively, and S is the stimulus presented. A graph showing the effects of the equation can be seen in figure 1.

3.2 Using Habituation in a Novelty Filter

In effect, habituation allows an animal to ignore stimuli that are seen often, so that the animal can concentrate on other, potentially more important features. This is exactly the functionality required of a novelty filter – removing stimuli that are seen frequently. The novelty filter that we describe uses habituation in conjunction with a clustering neural network to highlight novel stimuli.

The novelty filter uses a clustering network to learn a representation of the robot's perceptions. Input vectors of the robot's perceptions are presented to the clustering network, which finds the best-matching neuron using a winner-takes-all strategy. Each neuron in the map field is connected to the output neuron via an habituable synapse, so that the more frequently a neuron fires, the lower the efficacy of the synapse and hence the lower the strength of the output. The behaviour of the habituable synapses is controlled by equation 1. The strength of the winning synapse is taken as a novelty value for the particular winning neuron, and hence the perception presented, with more novel stimuli having values closer to 1, and more common stimuli values closer to 0.

We have experimented with a number of different clustering networks for use in the novelty filter. The one that performs best is a network designed especially for the task. The network, termed the 'Grow When Required' (GWR) network, has the capability of growing a new node whenever none of the nodes currently in the network matches the input sufficiently well. This means that the network grows until it is sufficiently large to represent the inputs to the required accuracy. Details of the GWR network and demonstrations of the novelty filter implemented with this network for sonar-based environment inspection are given in [8].

4 Image Processing

Images are captured from the camera as a two dimensional array of greyscale pixels. The image capture process is described in section 6. Before the novelty of each image can be evaluated, an input vector must be generated that is suitable for presentation to the novelty filter. However, the quality of the images was not high, and therefore a number of different preprocessing techniques were tried to improve the quality of the image. We report only those algorithms that were used. General references that deal with the types of image processing described are [3, 10].

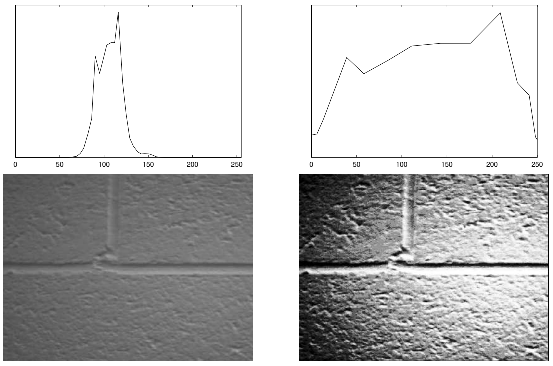

Figure 2: *Top: Left:* A typical histogram of image intensity before histogram equalisation. *Right:* The same histogram after histogram equalisation. *Bottom:* The images that produced the histograms.

4.1 Histogram Equalisation

The images were of low contrast. This can be seen by looking at the histogram of the brightness of the image over the range of greyscales (0 – 255). The histogram is sharply peaked, with very little spread of values, as can be seen on the top left of figure 2. Histogram equalisation flattens the peaks and enhances the minima of the histogram by increasing the distance between the points where the histogram is large, so spreading the histogram out to cover all the available image intensities.

5 Producing the Input Vectors

Once the image has been preprocessed, it can be used as input to the novelty filter. This requires the generation of an input vector. Generating the input vector reduces the amount of information that is stored, and therefore the choice of technique is very important. We tried several methods of producing input vectors from the images. In this section we describe the three that worked best.

The inputs to the novelty filter were all scaled to lie between 0 and 1. This was done during the calculation of the input vector, using the techniques described in the next section.

5.1 Fingerprints

The most simple technique, which also seems to perform the best on the novelty detection problem described in this paper, was to use as an input vector the intensity values of the image at particular points. After experimentation with a number of different shapes, it was found that using inputs from the spiral shown in figure 3, starting at the centre and working out to the edge produced the most useful information. This provides most information from the centre of the image, but does include pixels from the edges. It is this technique that was used in the experiments reported in this paper.

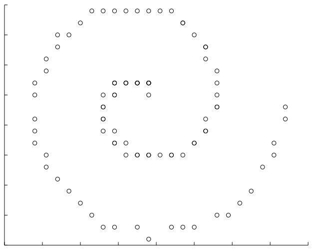

Figure 3: The spiral pattern used to generate the 'fingerprint' input vector to the novelty filter.

5.2 Histograms

Another technique that was tried was to take histograms in both the $x-$ and $y-$ directions. The histogram could be taken of the intensity values of the images directly, or an edge detection routine could be applied first, and then the histogram of the edge-detected image taken. It was found that using a simple edge-detection routine first improved the performance of this filter.

5.3 Principal Component Filters

We also tried an algorithm based on Sanger's image compression algorithm [9]. The image is partitioned into non-overlapping 8 x 8 squares of pixels. The Generalised Hebbian algorithm [9] is used to generate the first eight principal components of each small square of the image. The weights generating these principal components are saved and the trained network is applied to squares taken from each successive image. The principal component values produced by each of the squares is stored as a representation of that square. These principal component values are used as input to the novelty filter. If the first image if not representative of all the images, this filter does not work well, which makes it unsuitable for online use.

6 Experiments with the Novelty Filter

6.1 Experimental Procedure

The experiments described in this paper parallel those used to investigate novelty detection with sonar

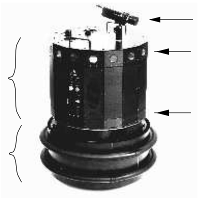

Figure 4: The Nomad 200 robot used in the experiments. Note the position of the camera at the top of the turret.

perceptions [7]. The Nomad 200 robot shown in figure 4 explored an environment from some arbitrary starting point by using a wall-following behaviour to travel adjacent to one wall, staying between 20 cm and 60 cm away from the wall. The robot stopped every 10 cm to take a photograph.

As can be seen in figure 4, the camera is on the top of the robot, about 1 m off the ground. In these experiments the camera faced the wall that was being followed. This meant that each image is of a small section of wall, as is shown in figure 2. The preprocessed photograph was presented to the novelty filter, which produced a novelty value. Before the experiment began, the camera was focussed manually. However, the focussing ring on the camera slipped over time, and the robot did not stay a constant distance from the wall, so some of the images are blurred.

Once the robot had travelled 10 m in an environment, the experiment was paused, and the robot was returned to the beginning of the run. After each pass through the environment with the novelty filter learning a model of the environment, a second pass was performed with learning turned off. The novelty filter still produced a value of the novelty for the current perception, but the network did not learn. This enables us to see how much was learnt during the previous run.

6.2 Results

Initially the novelty filter was initialised randomly, and the robot was allowed to explore an environment. The results of this exploration are shown on the left of figure 5, which is labelled environment A. The diagram at the top of the figure shows the appearance of the environment. The robot is shown facing in the direction of travel adjacent to the wall that it followed, and that the camera faced. The graphs show the amount of novelty in the image at that point, with a spike signifying total novelty and no spike meaning that the image is similar to previous perceptions.

Once the novelty filter had learnt an accurate model of environment A (meaning that no perceptions were found to be novel), which took four learning runs, the environment was changed. A door on the right-hand side of the robot, the side that the camera faced, was opened. This changed the perceptions at that point. The outputs of the novelty filter for these perceptions are shown on the right of figure 5, labelled environment A*. It can be seen that the perceptions of the doorway are found to be novel, but that no other part of the environment is. This is to be expected, as the rest of the environment is identical to the first one.

The novelty filter trained on environment A was also used in two different environments. Figure 6 shows the amount of novelty found when the novelty filter trained in environment A was placed into a new environment. The first environment (shown on the left of the figure) is environment A*, as shown in figure 5. The results for two new environments are also shown in figure 6. It can be seen that the amount of novelty found in environment A* after learning in environment A is low, as would be expected since they are the same section of corridor with only one change made.

The first of the new environments (environment B) is a similar section of corridor. The amount of novelty found is this environment is shown in the middle graph of figure 6. The novelty filter recognised many of the perceptions, but found the images of the doorways to be novel. There are two reasons for this. One is that the novelty filter has only seen perceptions of one doorway, so they are still seen to be slightly novel, and the other is that the doorjambs are painted black, instead of white as they were in the first corridor. This is why there is still a lot of novelty found in this environment, although less than was found in environment A with the untrained filter.

Environment C, the final environment, is a very different corridor. This can be seen from the fact that the novelty filter trained in environment A finds this environment very novel (shown on the right of figure 6). Instead of breezeblock on the wall, there are bricks, so the appearance of the wall is very different. The wall was found to be novel in the two runs in this environment, although the novelty filter soon learned to recognise it. In addition, at the end of the run are a number of posters. These were also found to be novel.

7 Conclusions

In this paper we have demonstrated that a novelty filter capable of learning online can be used on a mobile robot equipped with a camera to learn an internal model of an environment perceived through the images captured by the camera. Once this model has been acquired, the perceptions highlighted by the model are those that do not fit into the pattern of previous inputs.

Experiments demonstrating the effects of the filter

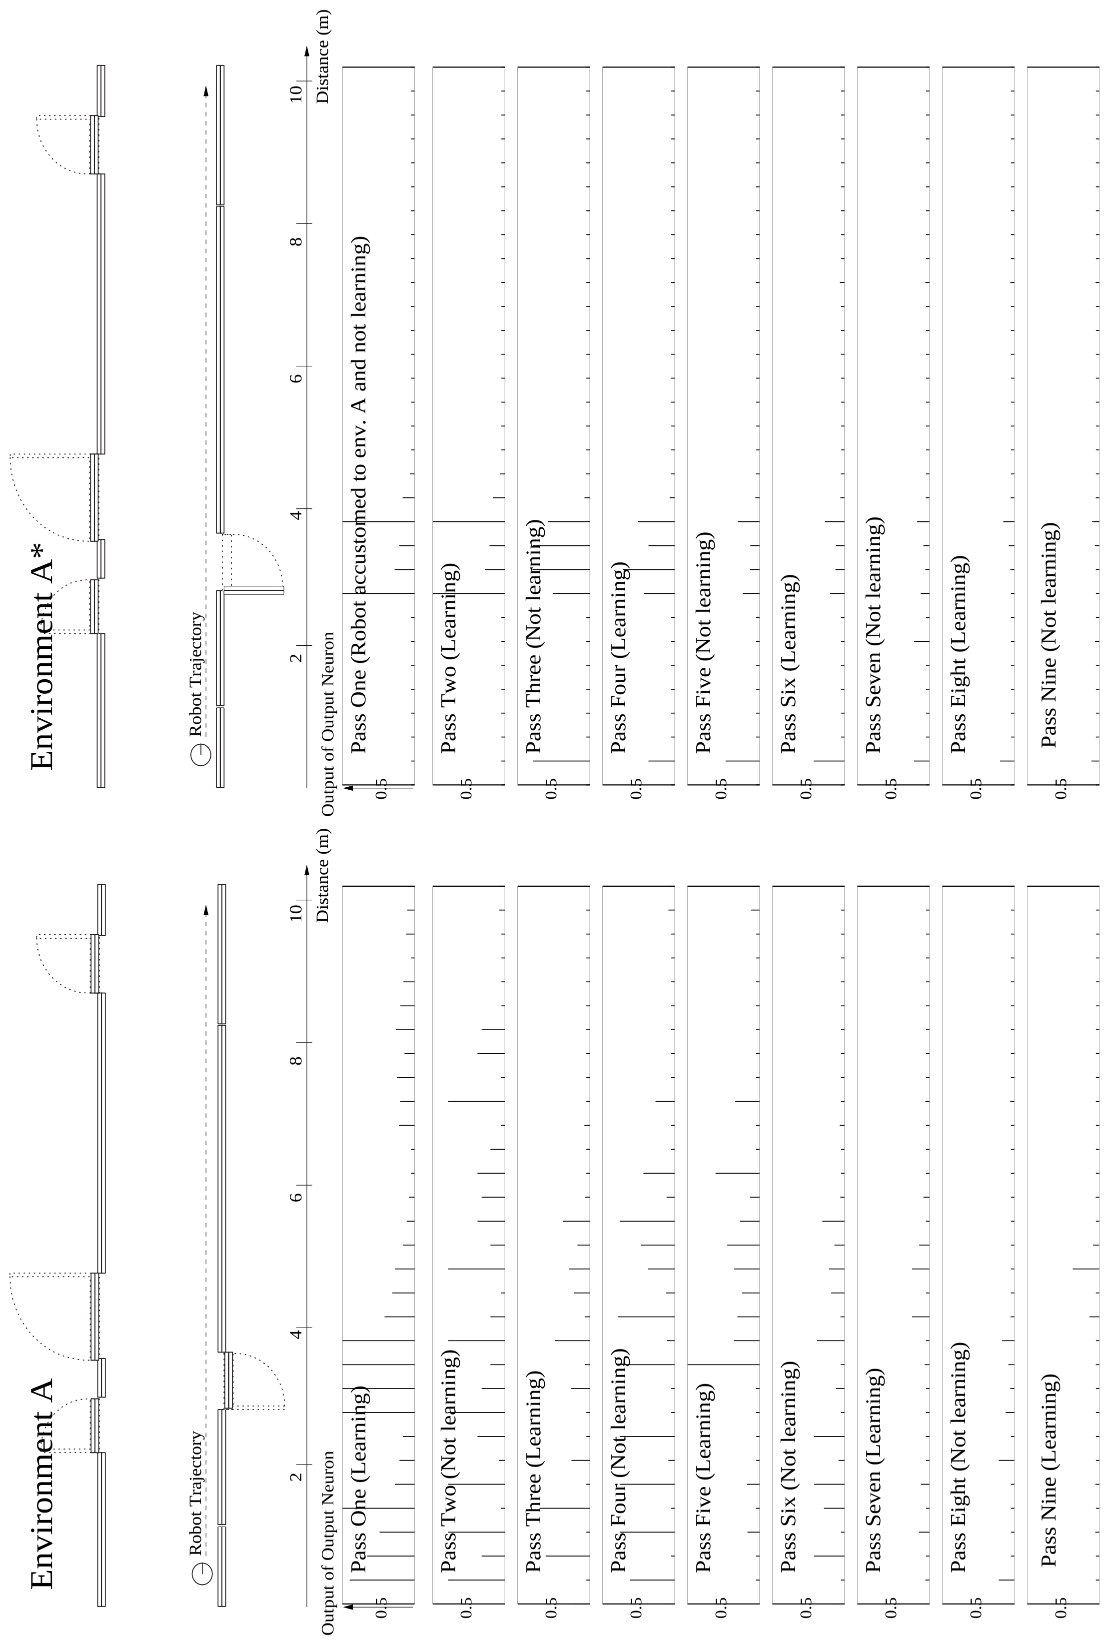

Figure 5: *Left:* The output of the novelty filter when the robot explores an environment without prior training (peaks are perceptions with high novelty values). Initially everything is novel, but the novelty filter soon learns to recognise perceptions such as the wall, that are seen often. *Right:* After training in environment A a door is opened. This changes the perceptions around the doorway and the robot explores the changed environment. It can be seen that only the changed perceptions of the doorway are found to be novel.

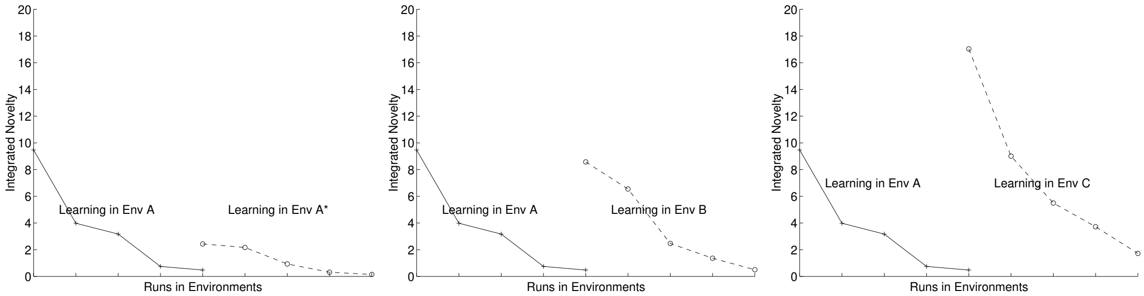

Figure 6: The results of the first inspection experiment. The graph on the left shows the robot learning first in environment A, and then in environment A*. It can be seen that once the robot has learnt about environment A, very little in A* is novel. The middle graph shows how the amount of novelty increases when the robot explores environment B after learning about environment A, and finally the graph on the right shows how the novelty increases when the robot begins to explore environment C.

on perceptions of a number of different environments that the robot explored are given. For reasons of space, the results have only shown the effects of one of the ways of generating an input vector for the novelty filter from the image and preprocessing has not been discussed fully.

One area that has not been addressed in this paper is that of sensor fusion. Previous experiments have demonstrated that the novelty filter works well with inputs from the robot's sonar sensors, and this paper demonstrates that the filter works with images from a camera. It is possible that using both types of sensor as input would also be useful, this is under investigation. Another area where work is needed is to investigate whether the novelty filter can be used to differentiate between different environments. This would enable robot behaviours to be tailored to the particular environment that the robot was currently in.

Acknowledgements

This research is supported by a UK EPSRC Studentship.

References

[1] Rafal Bogacz, Malcolm W. Brown, and Christophe Giraud-Carrier. High capacity neural networks for familiarity discrimination. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN'99)*, pages 773 – 778, 1999.

[2] Paul Crook. Spotting novelty: A neural network model for familiarity discrimination. Master's thesis, Division of Informatics, University of Edinburgh, 2000.

[3] Rafael C. Gonzalez and Paul Wintz. *Digital Image Processing.* Addison-Wesley, Reading, Massachusetts, 2nd edition, 1987.

[4] Teuvo Kohonen. *Self-Organization and Associative Memory, 3rd ed.* Springer, Berlin, 1993.

[5] Teuvo Kohonen and E. Oja. Fast adaptive formation of orthogonalizing filters and associative memory in recurrent networks of neuron-like elements. *Biological Cybernetics*, 25:85–95, 1976.

[6] Stephen Marsland, Ulrich Nehmzow, and Jonathan Shapiro. Novelty detection for robot neotaxis. In *Proceedings of the 2nd International Symposium on Neural Computation (NC'2000)*, pages 554 – 559, 2000.

[7] Stephen Marsland, Ulrich Nehmzow, and Jonathan Shapiro. Novelty detection on a mobile robot using habituation. In *From Animals to Animats: Proceedings of the 6th International Conference on Simulation of Adaptive Behaviour (SAB'00)*, pages 189 – 198. MIT Press, 2000.

[8] Stephen Marsland, Ulrich Nehmzow, and Jonathan Shapiro. A self-organising neural network that grows when required, 2001. Submitted to Neural Networks.

[9] T.D. Sanger. Optimal unsupervised learning in a single layer feedforward network. *Neural Networks*, 2:459 – 473, 1989.

[10] Milan Sonka, Vaclav Hlavac, and Roger Boyle. *Image Processing, Analysis and Machine Vision.* Chapman and Hall Computing, London, UK, 1993.

[11] James C. Stanley. Computer simulation of a model of habituation. *Nature*, 261:146–148, 1976.

[12] Odin Taylor and John MacIntyre. Adaptive local fusion systems for novelty detection and diagnostics in condition monitoring. In *SPIE International Symposium on Aerospace/Defense Sensing*, 1998.

[13] Alexander Ypma and Robert P. Duin. Novelty detection using self-organizing maps. In *Proceedings of International Conference on Neural Information Processing and Intelligent Information Systems (ICONIP'97)*, pages 1322 – 1325, 1997.