

Adding 'Foveal Vision' to Wilson's Animat*

Dave Cliff^{1,2} and Seth Bullock¹

¹School of Cognitive and Computing Sciences

²Neuroscience IRC, School of Biological Sciences
University of Sussex, BRIGHTON BN1 9QH, U.K.

davec or sethb, both @cogs.susx.ac.uk

Abstract

Different animals employ different strategies for sampling sensory data. The strategies are often closely constrained by environmental considerations, such as the animal's ecological niche. In animals that can see, differences in sampling strategy manifest themselves as differences in field of view and in spatially variant sampling (so-called "foveal" vision). In analysing adaptive behaviour in animals, or attempting to design autonomous robots, mechanisms for exploring variations in sensory sampling strategy will be required. This paper describes our work exploring a minimal system for investigating the effects of variations in patterns of sensory sampling. We have re-implemented Wilson's (1986) animat, and then experimented with altering its sensory sampling pattern (i.e. its *sensory field*). Empirical results are presented which demonstrate that alterations in the sensory field pattern can have a significant effect on the animat's observable behaviour.

Analysis of our results involves characterising the interaction between the animat's sensory field and the environment within which the animat resides. We found that the animat's observed behaviour can, at least in part, be explained as a result of the animat *cautiously* moving in a manner which maximises the uptake of *new* information from the environment over time.

1 Introduction

In the natural world, a vast range of *sensory sampling strategies* have evolved to enable animals to engage in adaptive behaviour. A sensory sampling strategy is a particular solution to the problem of extracting useful information from the environment, subject to certain constraints.

In either analysing the sensory sampling strategy of an animal, or devising a sensory sampling strategy for an artificial autonomous agent, there are a wide variety of issues to take into account. In both cases, intuitions and *a priori* notions may do more harm than good. If we are ever to succeed at either task, mechanisms for exploring the space of possible strategies will be of great value.

This paper discusses our recent work in studying the effects of varying the sensory sampling strategy in an elementary virtual autonomous agent: Wilson's animat (Wilson, 1986). The animat exists in a cellular world, and has (in comparison to a real animal) a small number of possible sensory inputs or behavioural

*© 1992 D.Cliff, S.Bullock. All rights reserved. Communication should be directed to D.Cliff.

outputs. The animat employs mechanisms, inspired by evolution and genetics, which lead to emergent adaptive behaviour. In its original formulation, Wilson's animat sampled its surroundings in a uniform manner. We have re-implemented Wilson's animat, and then studied the effects of altering its sampling pattern. We refer to the animat's sampling pattern as its *sensory field*.

For reasons of brevity, we assume that the reader is familiar with the fundamentals of Wilson's work: a brief overview is presented in Section 3.1.

We found that varying the sensory field caused significant differences in the animat's observable behaviour. These changes in behaviour were predictable from the hypothesis that for certain sensory field patterns, a strongly adaptive behaviour is to cautiously move in a fashion which maximises the amount of new information received over time.

The work described here is empirical and pretheoretical. We have chosen Wilson's animat primarily because it offers a convenient *minimal* system for study. We are acutely aware of the limitations of the path we have followed, especially of the highly idealised notions of perception and action involved.

2 Rationale

The rationale for this work arises from our current major research project at Sussex, working on evolving network control architectures for visually guided robots. One of the major obstacles in this work is computational cost. This is because we use accurate physics simulation models, and ray-tracing graphics for synthesizing visual input. Even with state-of-the-art computational resources, progress is relatively slow.

In order to gain more rapid access to some of the fundamental issues in sensory sampling strategies, we have pursued these studies using Wilson's animat: the minimal nature and relative simplicity of the system makes proportionately fewer computational demands, and thus allows for much more rapid progress.

Furthermore, the results from the animat system appear to be of significance in analysing biological systems: this is discussed further in (Cliff and Bullock, 1992).

Regardless of the simplified perceptual mechanisms involved in Wilson's animat, we still view this work as dealing with issues in visual perception. A significant influence in our approach to equipping robots with visual capabilities comes from studies of vision in insects. Typically, the compound eyes of insects form images using relatively few photoreceptors. Yet despite the low pixel-count, insects can exhibit complex and sophisticated visually guided adaptive behaviours. Often, this is made possible by use of *spatially variant* sampling. This is where the angular resolution of the pixels in the image is not uniform across the surface of the eye. Typically, there are a small number of narrow-angle high-acuity photoreceptors which sample a restricted area of the visual field at high resolution; while the remaining majority of the pixels sample at a lower resolution, but with a correspondingly wider acceptance angle. The animal's visual capabilities can thus be thought of as satisfying a particular trade-off between number of photoreceptors, visual resolution, and angular extent of the field of view.

Spatially variant vision is not particular to arthropods. In humans, it is more widely known as *foveal* vision; after the *fovea*, a small pit in the centre of the

retina where the high-acuity photoreceptors are concentrated.

Our aim then, was to alter Wilson's animat in ways inspired by foveal vision, and observe the results. We could, in principle, have devised our own animat system, perhaps with a more complex environment, or operating with a different adaptive algorithm. However, the great advantage of Wilson's system is that it is well documented, and has been replicated by other authors. If we had developed our own animat, there would always have been the hidden danger that some unforeseen interaction in the system could interfere with our interpretation of the results. Wilson's animat is a sufficiently mature piece of research that we felt this was a much less likely event.

3 Experiments

3.1 Background

Wilson's animat exists in a regular rectangular grid of cells. It occupies one cell, and can move from cell to cell within the grid. In our work discussed here, we have used a particular cellular environment, 58 columns by 18 rows, referred to as woods7: see Figure 1. This environment was used by Wilson in his work (Wilson, 1986, p.260). Each cell may either be empty, or contain food (represented as F in Figure 1), or an obstacle (represented as T for 'tree'). The position of the animat is represented by an asterisk (*). The animat can move by making a 'step' into any one of the eight cells surrounding it, but it cannot step into a cell occupied by a T. If the animat steps on a F, it receives a reward of fixed amount (the *payoff*). The 2D world is a toroidal surface. The next few paragraphs summarise the details of Wilson's work which are cogent to this paper: in the interests of brevity, many important details are omitted, and the reader is referred to (Wilson, 1986) for further information.




Figure 1: The environment "woods7". After (Wilson, 1986, p.260). The compass is arbitrary, but is used in the text.

The animat's initial position is a randomly chosen empty cell. At each position, the animat samples a sense-vector. The sense-vector is formed by examining the eight cells surrounding the animat: see Figure 2. It is this sampling of eight nearby cells which we refer to as the animat's sensory field.




Figure 2: Formation of the sensory vector and its related binary encoding. Left figure shows numbering of the animat's surrounding cells (arbitrary, but consistent); animat represented by *. Centre figure shows detail of animat in its environment: 3 trees and 2 foods are present, but the animat can only 'see' one tree and one food. Table shows correspondence between cell-number and content (T=tree; F=food; E=empty)

The animat's control system is a large population of simple condition-action rules, called classifiers. The condition part of the classifier is a template against which the sense-vector is compared. For a given sense-vector, a list is compiled of all matching classifiers in the population. Each classifier has a strength value associated with it: initially all classifiers have the same strength, but over time the strength of each classifier alters to reflect its usefulness in finding food.

When the animat receives payoff for stepping onto food, the reward is distributed among the recently matching classifiers, increasing their strengths. After updating these values, there is a possibility that the classifier population will undergo some genetically-inspired alteration. Finally, the animat is randomly repositioned to a blank cell somewhere in the environment, and the search for food recommences. The animat's performance is rated on the basis of the number of steps it takes to find food.

The animat's initial population of classifiers is randomly generated, and its initial performance is normally no better than random search, but over many food-finding cycles, the performance improves significantly.

The rate at which performance improves, and the final average performance value arrived at, is dependent on fine-tuning a number of parameters: see (Cliff and Bullock, 1992) for further details. Using *woods7*, we were able to qualitatively reproduce Wilson's published results. We were not concerned with quantitative reproduction, as the intention in this work is not necessarily to *improve* on the performance of Wilson's system: we are primarily interested simply in recording any variations in the performance of our implementation brought about by changes in sensory field. The following sections describe our approach.

3.2 Approach

This section gives details of our methods, and the techniques we use to present our results. Methods are described in Section 3.2.1; Section 3.2.2 discusses analysing the environment, an essential move in interpreting subsequent results.

3.2.1 Methods

Before introducing variations in the sensory field, we ran control experiments on the animat with the sensory field used by Wilson. In his publications, Wilson calls his animat "***"; to distinguish it from animats with other sensory fields, it is referred to as **w* from now on. Our aim in performing control experiments was to

elucidate the 'normal' pattern of behaviour of *w, in order that we could identify any significant variations arising as a result of altering the sensory field.

The mechanisms underlying adaptation in the animat are nondeterministic, and so it was necessary to observe the behaviour of *w over a large number of trials, and then calculate summary statistics. We tested *w a number of times, where each test involved the animat finding 8000 items of food. On each 8000-food test, our aim was to monitor the behaviour of the animat once it had sufficiently 'matured' such that it was performing efficient food-finding behaviour. We did this by recording the movements of the animat over the last 2000 food-finding trials in each 8000-food test (i.e. the last 25% of its 'lifetime'). Typically, the animat had developed a 'preference' for consistently wandering in one or two of the eight possible directions, until a T or F came into 'view' (i.e. until the sense-vector was not entirely empty). Once a T or F came into view, the animat would typically manoeuvre in order to move onto the food.

Because of this emergent 'preference', we repeated the 8000-food test 1200 times, starting each time with a 'fresh' animat (i.e. one with a randomly initialised classifier system). For each of the 1200 animats with the *w sensory field, we recorded the moves made over the last 2000 food-finding tasks.

Viewing the results from 1200 *w animats, we found that the 'preferred' direction varied from animat to animat. However, our primary interest is at the population level (i.e. the expected 'average' behaviour of the mature animat with a particular sensory field). To this end, we present here only the average behaviour witnessed in the 1200×2000 monitored food-trials.¹ We calculated average behaviour by counting how many of the monitored movements were made in each of the eight directions, and then dividing the count for each direction by the total number of moves recorded. This tells us the proportion of moves in each direction made by the 1200 'mature' animats. The results for *w are illustrated in Figure 3.




Figure 3: Illustration of the proportion of moves made by *w in each of the eight directions over 1200×2000 food-finding tasks. The length of the vector in each direction indicates the number of moves made in that direction (as a proportion of the total number of monitored moves). The octagon is provided for reference: if the proportion of movements in each direction was identical, the endpoints of each vector would touch the edge of the octagon.

It is this statistic that forms the basis of comparison in our analysis of the effects of varying the sensory field in manners inspired by foveal vision. However, in determining the average behaviour of a population of animats, we found that explanation was not possible without some characterisation of the environment in which the animats were situated. The next section deals with analysing the environment.

¹Note that the number of steps taken to find food varies, depending on the randomly-chosen starting-point, the random initialisation of the classifier system, and the particular sequence of sensory inputs received.

3.2.2 Analysis of the Environment

Examination of the woods7 environment indicates that, although it is not highly ordered, there are some observable regularities. The distance from any F cell to its nearest neighbouring F cell is fairly constant, and each F cell has two T cells immediately adjacent. Such environmental regularities will be exploited by the animat's adaptation mechanisms: for example, if a T cell is within the sensory field, then food cannot be very far away; therefore it will probably pay the animat not to move *away* from T cells, but rather to manoeuvre *around* the T cell until the F cell becomes 'visible'. If the environment was totally devoid of structure or regularities, then the animat could do little better than random search.

Because the animat is likely to exploit these observed regularities, it is important to analyse the environment. This is of most use in interpreting the average behaviours of animats (where "average" is used in the sense defined above). One might naively expect that the movements of mature *w animats should on average show no bias for any particular direction, given that there is no directional bias in the sensory field. But this is not the case.

Rather, as can be seen from Figure 3, we found non-uniformities in the distribution of recorded moves (e.g. the proportion of north/south movements was smaller than the proportion of movements in other directions). This may not simply be explained as sampling error. Instead, it can be explained by considering results that would be expected of an 'ideal' animat, which always moves in the shortest path towards F cells.

To gain insight on the expected ideal behaviour, we analysed woods7 by calculating the minimum number of moves required to find food from each cell in the environment. For details and illustration of this process, see (Cliff and Bullock, 1992).

Having derived the distance-to-food map, we examined each empty cell, and determined in which direction(s) a step would *reduce* the distance to food (in many cells there is more than one such direction). For each direction that reduces food-distance, we incremented a counter. When each cell had been processed in this fashion, we divided the count for each direction by the sum of all the counts, giving a proportion of best-possible-moves score for each direction. The relative proportions are shown in Figure 4. As can be seen, the population average for an ideal mature animat would show a bias for diagonal movements, with comparatively few north-south movements expected. Comparison of this 'ideal' with the results from *w (Figure 3) indicates a qualitative similarity (e.g. in rank ordering).




Figure 4: Illustration of the proportion of moves made by 'ideal' mature animat with omniscient sensory system.

However, *w could never approach the ideal because when it is in 'free-space' (i.e. when the sense-vector is all empty), there is no information to guide it to food. In such situations, wandering consistently in a certain direction (with occasional

perturbations to avoid periodic motion on the toroidal surface of the environment) until the sense vector is not all-empty is a sensible behaviour, in comparison to e.g. on each step moving in a random direction and relying on "drunkard's walk" phenomena. This explains both the emergent 'preferences' discussed above, and also dooms the animat to suboptimal search. Because the search in free-space is suboptimal, the average behaviour of mature *w cannot be expected to match the expected ideal average behaviour too closely.

Having satisfied ourselves with an understanding of the observed average behaviour of mature *w, we set about altering the sensory field and recording any resultant deviations in average mature behaviour.

3.3 Altering the Sensory Fields

In altering the sensory field (SF hereafter), we imposed constraints in order to limit the space of possible SF's considered. The constraints we employed were somewhat *ad hoc*, but did have some supporting rationale from considerations of 'real' (biological) vision. First, we constrained all the SF's to have exactly 8 cells, the same number as for *w (increasing or decreasing the number of cells severely complicates any comparisons as the dimensionality of the animat's sensory space is altered). Second, we ensured that the altered SF's were bilaterally symmetric about the north-south line through the animat's cell. Third, we constrained the layout of the SF's so that each cell was connected to either the animat or another SF-cell in one of the directions north, south, east, or west. Finally, we arbitrarily biased the alterations to produce "south-heavy" SF's: we required that the alterations moved SF-cells from the north half to the south half of the animat's surrounding cells; this halved the space of possible SF's we had to consider, because it eliminated "north-heavy" fields which were merely reflections of "south-heavy" ones. This had the effect of introducing "blind-spots" on the north side of the animat. Nevertheless, the toroidal nature of the animat's world allowed the animat, in principle, to always move south (where its 'vision' is better), and hence to avoid moving in a direction in which it is blind.

We have experimented with a number of different SF designs, each of which was labelled with an arbitrarily chosen identifying letter. The sensory fields we investigated are illustrated (along with experimental results) in Figures 5 to 10.

These SF's were selected because they are all analogous, in a sense, to feasible visual sampling strategies. Wilson's animat *w can only detect objects of interest in its immediate surroundings, whereas animats *s, *t, and *v can sense objects at a greater distance, but have suffered a corresponding loss in the angular extent of their sensory fields. Metaphorically, we can think of *w as being equipped with 'wide-angle lenses', while *s, *t, and *v have 'telephoto lenses'. While *w might be considered as trying to maximise the coverage of its field (at the expense of its degree of look-ahead), *v is maximising its look-ahead, but as a consequence has very narrow coverage - it can 'see' well in only one direction. The other animats can be thought of as, to greater and lesser degrees, exhibiting trade-offs between depth and breadth of coverage.

Having devised these new sensory fields, we set about investigating their effects on the animat's behaviour. Each sensory field was tested by observing epochs of 1200 animats (each in their own solitary environment) for every SF pattern. Each

animat was tested over 8000 trials (i.e. each SF was tested in 1200 animats \times 8000 food-finding tasks). Results from mature animats using these sensory fields are presented in the following section.

3.4 Results

In Figures 5 to 10, the results from testing are shown alongside the sensory field for the animat. The results are presented using the format introduced in Figure 3. Discussion of these results follows in Section 3.5.



Figure 5: (Left) SF for *a; (right) results.



Figure 6: (Left) SF for *b; (right) results.



Figure 7: (Left) SF for *g; (right) results.



Figure 8: (Left) SF for *s; (right) results.



Figure 9: (Left) SF for *t; (right) results.




Figure 10: (Left) SF for *v; (right) results.

3.5 Discussion

As can be seen from the figures, the observed movements for all the mature animats with altered SF's show significant variations from the results for *w.

The severity of the variation (and to some extent the *nature* of the variation) appears to be related to the 'severity' of the alteration in SF (i.e. the degree of southwards bias introduced). Nevertheless, even *a (whose SF differs from *w's SF by only one cell) clearly shows variations in movement patterns: more movements with a southern component are made, with a corresponding reduction in northwards movements.

Clearly, the mechanisms employed in the animat have served to ensure that the animat's observable behaviour is adaptive with respect to the interactions between the animat, its SF, and the environment. Put most simply, this demonstrates that it is an adaptive behaviour for an agent to move in the direction where it has most sensory information: if *a moves in a direction where sensory information is poor (e.g. due north, where it has a 'blind spot'), there is the danger that it will hit an unseen T obstacle.

The reduction in proportion of moves with a northward component is in sharp contrast to the 'ideal' animat behaviour shown in Figure 4, where moves northeast and northwest are proportionately higher than moves in directions west, east, or south. Indeed, in all the non-*w animats, the proportions of northwest and northeast movements are reduced, which indicates a hypothesis that: *it is an adaptive behaviour for the animat to move in the direction where it has most sensory information, so as to avoid unseen hazards lurking in areas of poor or zero sensory information.*

However, this hypothesis does not completely account for the observed behaviours (e.g. *v should only ever move south): it doesn't take into account the number of "new" cells sampled by the SF for each direction of movement. By "new" cells we mean those cells which come within the SF after a step is made, that were not sampled by the SF before the step was made. This does not include a previously 'unseen' cell moved onto by the animat itself (because the animat can only act on the basis of cells sampled by the SF). For illustration of the number of 'new' cells uncovered by each animat for a step in each direction, see (Cliff and Bullock, 1992).

So perhaps, when in 'free-space', a good movement policy is not to move in the direction where most sensory data is currently available, but *to move in a fashion which maximises the uptake of new information from the environment on each timestep.* Naturally, once an F falls within the SF, this policy should be abandoned in favour of the animat taking a direct path towards the F.

This is well illustrated by the following *gedanken* experiment. If *v always moved south in free-space (with e.g. occasional or random west or east movements to break cyclic paths on the toroid), it only uncovers one 'new' cell on each step. It is better for *v to drift in directions with either an east or a west component, until an F comes within the SF. At that point it should travel south in order to move onto the food. If we choose to use intentional language, we could say that in the east/west movements the animat is "seeking" food, while in the southwards movements the animat is "approaching" food.

There is, however, at least one unresolved issue in this version of events: For *v, moving in any direction other than south is risky, in that the animat will be moving onto an 'unseen' cell. If it moves north, it samples no new sensory information *and* runs the risk of colliding with a T. However, for the other six directions it has a risk of collision but gains the advantage of sampling 8 'new' cells, and the risk is the same for each of the six directions. Thus we would expect populations of mature *v to move south frequently (when approaching food), and to drift equally in all other directions except north. This is a fairly close qualitative description of the experimental observations (Figure 10). Similarly, for the other animats with altered SF's, drifting south-east or south-west is preferable to drifting north-east or north-west, because the northern-diagonal movements run the risk of an

"unseen collision" while the southern-diagonal movements move the animat onto a 'safe' cell sampled by the SF before the step was made. Again, this analysis appears to account at the qualitative level for the experimental observations.

This issue of 'cautious' drifting seems to be the final point necessary to explain the experimental results.

3.6 Summary

We have demonstrated that the adaptive mechanisms in Wilson's animat can lead to differing observable behaviours in animats with different SF's. This leads to the hypothesis that patterns of adaptive behaviour are dependent on the interaction of the animat's SF with a particular environment, and that observed patterns of behaviour can be explained as an interaction between three phenomena: *Seeking*, where the animat moves in a manner that maximises the receipt of new information from the environment over time, attempting to identify areas of the environment where a payoff (F) may be found; *Approaching*, where the animat moves towards areas of the environment likely to provide payoff; and *Caution*, where the animat is more likely to make "seeking" movements in directions which *do not* involve the animat moving onto cells that were previously unsampled, because of the risks of 'unseen hazards'.

4 Conclusions

Our work has explored various aspects of the effects of varying the sensory field profile of Wilson's animat. We have demonstrated that the variations have a significant effect on the animat's external observable behaviour. The simplifying assumptions and minimalism of the approach have been vindicated, in that we have identified important factors governing the interplay between SF, environment, and behaviour, which seem to be true of both animats *and* animals (discussed further in (Cliff and Bullock, 1993)). However, we suspect that complexity barriers are likely to be encountered if this cellular-world approach is extended: for truly general results, the noisy, continuous, dynamical, uncertain, and hostile nature of real worlds should not be ignored any longer.

Acknowledgements

Thanks to Stewart Wilson for his comments, and Phil Husbands for use of his exquisite C code.

References

- [1] D. T. Cliff and S. G. Bullock. Adding 'foveal vision' to Wilson's animat. Technical Report CSRP 263, University of Sussex School of Cognitive and Computing Sciences, December 1992.
- [2] S. W. Wilson. Knowledge growth in an artificial animal. In K. S. Narendra, editor, *Adaptive and Learning Systems*, pages 255-264. Plenum Publishing Corporation, 1986.