

Adaptation of Obstacle-Avoidance in the face of Emerging Environmental Dynamics

Chris Thornton
Cognitive and Computing Sciences
University of Sussex
Brighton BN1 9QN
Email: Chris.Thornton@cogs.susx.ac.uk
Tel: (44)273 606755 x 3239

June 29, 1994

Abstract

The avoidance of static obstacles is relatively easy to implement as an animat behaviour, even where the animat receives realistic ‘sensory’ input from the environment (e.g., simulated, infra-red proximity signals). [1,2,3] The behaviour can also be learned by a variety of methods. However, natural implementations of static-obstacle avoidance may break down if the objects in the environment have dynamic properties. Any introduction of object motion directly affects the environmental cues which anticipate impending collisions. The paper shows that this happens in a systematic way. The statistical order of the sensory collision cues increases roughly in proportion with the level of obstacle mobility. Thus, adaptation of obstacle-avoidance behaviours (in animats) involves making a transition from strategies based on low-order cues to strategies based on high-order cues.

1 Introduction

The aim of the paper is to investigate how simple, animat-based implementations of static-obstacle avoidance can be adapted to emerging environmental dynamics (i.e., the increasing mobility of environmental objects). First, the performance of a non-adaptive, obstacle-avoidance strategy is investigated using a 2-dimensional, animat simulation environment.¹ The simulations show that though the implementation of the behaviour is completely successful under static environmental conditions, it is quickly defeated once environmental dynamism increases beyond a certain threshold. I analyse the statistical order of the environmental cues upon which (successful production of) the behaviour is based and show that the initial implementation is sensitive to first-order cues only. As the environmental dynamics increase, the order of the significant environmental cues increases and this renders the implementation ineffective.

2 The simulation setup

A conventional animat simulation was used for the experiments. The simulated world was a 2-dimensional, rectangular space containing six circular obstacles. The animat was triangular in shape and was capable of making forwards, backwards and rotational moves within the space. At the beginning of each simulation it was placed into the world at a central position and on a random heading (i.e., facing in a randomly chosen direction). Its behaviour in each simulation cycle was produced by a hand-crafted

¹The environment is provided by the ‘POPBUGS’ library, running under the Poplog environment. Poplog users can obtain this software free of charge from the author.

controller which generated levels of ‘drive’ to be applied to two latitudinally-mounted wheels. These drive levels were interpreted within the simulation program in the obvious way. The animat’s movement in each cycle included a rotational component derived from their difference and a forwards/backwards component derived from their sum.

The animat was equipped with 7 proximity detectors. In each simulation cycle, each proximity detector produced a value representing the proximity of the nearest object along a particular ray. The rays were arranged in a forwards-facing fan, with one ray pointing directly ahead and each ray offset 15 degrees from its neighbour(s). The proximity values were normalised with respect to the maximum measurable distance within the space. Thus proximity values close to 1 indicated the presence of very near obstacles while the proximity value 0 indicates no obstacle detected.

The general form of the simulation setup is shown in Figure 1. Here we see the animat shown as a triangle in the centre of the space. The six circles represent the six obstacles. The dashed lines show the proximity-detector rays and the small integers show the points at which the rays have intersected with an obstacle or a boundary of the space.




Figure 1:

3 Static-obstacle avoidance

The animat’s control procedure was based on the the following two rules.

- If one of the proximity inputs exceeds 0.9 then execute a small turn to the right.
- Otherwise execute a small forwards move.

Under this control regime, the animat ‘wanders’ through the space, turning away (to the right) whenever it is confronted with an obstacle or a boundary. The general effect is one of a random, obstacle-avoiding exploration of the space; see Figure 2.

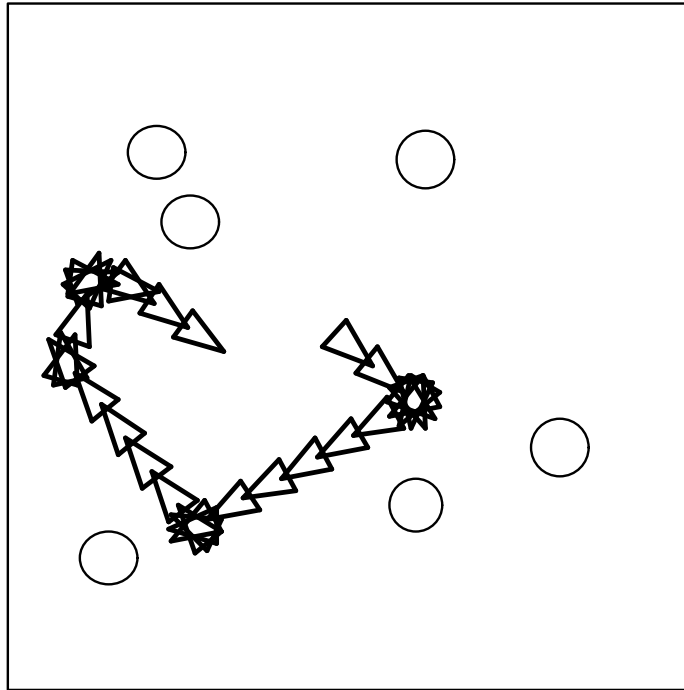


Figure 2:

4 The introduction of environmental dynamics

The impact of emerging dynamics was examined by modifying the simulation so as to give the obstacles simple dynamic properties. Each obstacle was randomly assigned a heading and then moved forward a certain distance in each simulation cycle. An obstacle arriving at a boundary would have its heading increased by 170 degrees (modulus 360), i.e., nearly reversed. This ensured (1) that obstacles never moved beyond the boundaries of the space and (2) that they would not simply track back and forth on the same linear trajectory.

A ‘dynamism’ parameter was implemented to control the degree of mobility of the obstacles. A dynamism value of 0 ensured that all obstacles would remain completely static. A dynamism value of 0.5 ensured that obstacles would have a limited mobility. A dynamism value of 1 ensured that obstacles would have a maximum degree of mobility, i.e., would move through the space very rapidly.

Using the modified simulation, it was easy to show that the original implementation of obstacle-avoidance was robust for low dynamism values. For example, consider Figure 3. This shows a simulation sequence in which the animat (using the control procedure described above) successfully negotiates a number of relatively slow-moving obstacles. The original implementation of the behaviour remains effective under this modestly dynamic regime largely because the animat is moving faster than the obstacles. The slowness of the dynamic obstacles means that they are reasonably well approximated as static objects and this allows the original controller to remain effective.

If we increase the level of environmental dynamism the effectiveness of the control procedure is significantly reduced. This effect is shown in Figure 4. Here the animat ‘crashes’ into a relatively fast-moving




Figure 3:

oncoming obstacle in the initial stages of the simulation.

5 Statistical analysis

The adaptation of even very simple behaviours (such as obstacle avoidance) to emerging dynamics is a *hard* problem and this paper does not, in any sense, attempt to provide a direct solution to it. What it will do is show that any successful adaptation strategy for static-obstacle avoidance must have certain properties. In particular the adaptation strategy must enable the controller to test environmental cues of increasing statistical order [4]. This is most easily demonstrated via an examination of the sensory inputs used by the obstacle-avoiding animat within the simulation.

The rows of Table 1 show a sequence of 10 sensory-motor associations implemented by the animat controller during the simulation. In each row, the first seven columns contain the proximity values produced in a particular cycle, and the final two columns show the drive levels returned by the controller procedure. Note that in the first three cycles shown, none of the proximity values exceeds the threshold value of 0.9. The controller responds by specifying full drive for both wheels (i.e., a small forwards moves). In cycles 4-6 the sensory inputs contain above-threshold proximity values. The controller responds by specifying left-wheel drive only (i.e., a small turn to the right). In the remaining cycles the proximity values are always below threshold and pure forwards moves are thus specified.

The controller used for the simulations implements a rule which refers to any *one* of the proximity inputs. This means that it relies on the testing of a first-order environmental cue (i.e., a sensory property involving a single variable). Provided obstacles are static, or nearly so, this approach presents no disadvantage. A first-order cue is capable of showing quite unambiguously that there is an object directly ahead. Thus, testing a first-order cue is sufficient to reliably determine whether a turn is necessitated.

However, consider what happens as we increase the level of environmental dynamism. As long as obstacles




Figure 4:

move relatively slowly, first-order cues are sufficient to determine the presence of an obstacle directly ahead. As the velocity of obstacles increases, first-order cues cease to be effective indicators of obstacle presence. A fast-moving object that is directly in front of the animat in one cycle of the simulation may move out of its path in the next. Alternatively — and more disastrously for the controller — a fast-moving obstacle that is well clear of the animat’s trajectory in one cycle may move inside it in the next. A first-order cue (i.e., a single proximity value) cannot discriminate between these two cases. Therefore single proximity values are necessarily ambiguous with respect to the task in hand.

With modest dynamics, then, unambiguous evidence concerning the presence of an obstacle in the region of space about to be occupied by the animat can only be obtained through tests made on more than one proximity value. In other words, under modestly dynamic conditions, environment cues become at least second-order. But we clearly need to determine the rapidity — relative to the increase in environmental dynamism — of the increase in the order of environmental cues.

0.792	0.799	0.791	0.767	0.719	0.599	0.309	1	1
0.739	0.747	0.738	0.707	0.64	0.495	0.28	1	1
0.683	0.695	0.685	0.647	0.567	0.391	0.248	1	1
0.937	0.935	0.688	0.694	0.68	0.637	0.546	0.3	0
0.932	0.936	0.937	0.94	0.69	0.694	0.673	0.3	0
0.744	0.925	0.932	0.936	0.944	0.931	0.693	0.3	0
0.697	0.696	0.868	0.876	0.876	0.63	0.655	1	1
0.649	0.648	0.805	0.816	0.524	0.589	0.618	1	1
0.602	0.601	0.571	0.756	0.478	0.552	0.58	1	1
0.554	0.553	0.521	0.696	0.432	0.511	0.544	1	1

Table 1:

Given the potential complexity of the effect there would seem to be little hope of finding an analytic solution. However, there are good reasons to think that cue order will increase roughly in proportion with environmental dynamism. In the simulations performed all the obstacles in the environment moved on linear trajectories at the same constant speed. This is clearly a minimally dynamic regime in which second-order environmental cues spanning no more than two sensory cycles can be used to unambiguously determine the linear trajectory of an object. This information is sufficient to determine the presence of the obstacle in the next simulation cycle. Thus, under the regime simulated, second-order, two-cycle environmental cues should contain sufficient information for effective control.

A modest increase in the qualitative character of the environmental dynamics can be obtained by allowing the obstacles to move on circular rather than linear trajectories. Under this regime, it is clear that at least three proximity measurements (spanning three cycles) are required to establish the trajectory of an obstacle. Environment cues for the behaviour thus become at least third order. If we now allow obstacles to adopt arbitrary dynamic characteristics (e.g., we allow obstacles to choose whether to move in a zig-zag motion, or a looping motion, or jerky motion) then we can expect the order of environmental cues for the behaviour to increase rapidly. At a rough estimate, an environmental cue for successful obstacle avoidance with obstacles producing arbitrary dynamics might be greater than tenth-order.

The general implication is that any successful adaptation of static-obstacle avoidance (and probably many other simple animat behaviours) which seeks to deal with emerging environmental dynamics must take account of the necessity to test increasingly high-order environmental cues spanning increasingly large numbers of sensory cycles. The observation is but a small step on the way to an effective decomposition of the *general* problem of adaptation to dynamic effects. But it does provide us with a rule of thumb for evaluating potential adaptation models. For example, imagine that the simple recurrent neural network (SRN) [5] is put up as a candidate model for the adaptation-to-dynamics task. The rule-of-thumb regarding environmental cues might lead us to attribute low plausibility to this model on the grounds that SRNs have a limited capacity for preserving persistent internal state and thus could not be expected to deal with high-order cues spanning many cycles of sensing and behaviour.

6 Summary

The paper has described some animat simulations which sought to investigate the preservation of obstacle avoidance behaviour in an increasingly dynamic environment. The simulations showed that an implementation which is highly effective in a static environment degrades rapidly as the velocity of the obstacles increases. Informal arguments were put forward in support of the view that there will typically be a rough proportionality between the increase in dynamism and the increase in the order of environmental cues which must be tested. This observation provides the beginnings of a decomposition for the problem of adaptation-to-dynamics and a potential rule-of-thumb for evaluating existing models.

References

- [1] Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. London: The MIT Press.
- [2] Kontoravdis, D., Likas, A. and Stafylopatis, A. (1992). Collision-free movement of an autonomous vehicle using reinforcement learning. *Proceedings of the Tenth European Conference on Artificial Intelligence*. John Wiley Sons.
- [3] Thornton, C. (1993). Supervised learning of conditional approach: a case study. CSRP 308, Cognitive and Computing Sciences, University of Sussex, UK.
- [4] Duda, R. and Hart, P. (1973). *Pattern Classification and Scene Analysis*. New York: Wiley.
- [5] Elman, J. (1990). Finding structure in time. *Cognitive Science*, 14 (pp. 179-211).