

Nature Inspired Target Reacquisition Strategies for Autonomous Vehicles in a Search Role

Alec Banks and Jonathan Vincent

(*Software Systems Research Centre Bournemouth University*)

(Email: al.banks@ntlworld.com)

Previous research by the authors has investigated the application of nature inspired search strategies and particle swarm optimisation to the guidance of autonomous vehicles in search roles. Findings suggest that nature inspired search strategies are beneficial under a wide range of search conditions. However, that research did not consider the potential for reacquisition of lost targets, which is especially of interest where the targets are travelling faster than the searching agents or in environments that allow for some means of visual or other sensory occlusion. This work investigates whether behaviours observed in nature can be applied to scenarios where target detection is lost. Through employing a variety of reacquisition strategies against a range of target movement types it is demonstrated that the overall effectiveness of the system can be improved in a variety of scenarios by incorporating nature inspired reacquisition strategies.

KEY WORDS

1. Natural search strategies.
2. Target Reacquisition.

1. INTRODUCTION. The authors' previous work, which provided an initial investigation into the application of particle swarms to the guidance of autonomous aerial vehicles (Banks et al., 2008), identified two main areas for further development to enhance the performance of swarms of vehicles defending an area against intruders. Both of these involved search activities: the first being the approach to initially locating targets and the second being the swarm behaviour when a target escaped detection, whereupon the existence of the target was immediately 'forgotten', leading to the defending vehicles resuming their initial search behaviour. The former of these was investigated and improved by exploring a range of nature inspired search strategies (Banks and Vincent, 2009), whilst the latter is investigated here.

Defending vehicles that forget targets once detection is lost are potentially undesirable since the information gleaned about the target during detection might be used to improve the likelihood of re-detection. Nature has evolved strategies that utilise the information about the last known location and behaviour of lost prey items to improve the chances of reacquisition. Such behaviours have developed because they can improve overall foraging efficiency since, if a prey item is lost, the initial effort in locating and pursuing it is wasted.

Although not exhaustive, this work examines the effect of employing a variety of reacquisition search strategies identified from existing ethological literature: predictive tracking, which uses the last known speed and direction of a target to predict its current location; direct search that assumes a target may have either hidden at its last known location or not moved far from it; saccadic movement, as observed in the eye movements of animals tracking a target that is faster than it can track due to limits on tracking ability caused by physiological limits and cognitive processing delays; and crosswind casting, which is a strategy employed by animals following (and losing) a scent trail. Each of the strategies is employed against a variety of target speeds and target movement types, as previously employed by Banks and Vincent (2009). Throughout the remainder of the work the term agents is employed to refer to the autonomous vehicles in recognition that the findings could be useful across a range of search applications, not just autonomous vehicles, where an entity can sense its environment and calculate a target's location, speed and direction.

Section 2 briefly reviews the ethological literature that underpins the implemented reacquisition strategies, which are described in detail in Section 3. Section 4 describes the environment and occlusion mechanism used to necessitate reacquisition. Section 5 presents the empirical work and results. Section 6 discusses the findings and Section 7 concludes.

2. BACKGROUND. This research builds on previous work (Banks *et al.*, 2008; Banks and Vincent, 2009) which should be referred to for more details regarding the swarming and initial search elements of the system along with more complete descriptions of the simulation environment and target behaviour. The remainder of the section reports existing ethological research that provided the inspiration for the target reacquisition strategies applied in this research followed by a brief section that touches on how a predator knows when to give up a particular search.

2.1. Reacquisition Search. Whilst in the pursuit phase of hunting, an animal may find that the sensory information about a target becomes obscured. At that point the predator may decide to give up or to apply a reacquisition strategy. There are many components to target tracking and to fully review them is beyond the scope of this work. However, some elements have been identified as having potential as metaphors for improving the systemic response to loss of target detection.

Direct search is a strategy whereby a predator will return to the last known location of a lost prey in cases of occlusion because the prey may have stopped to hide. Ogden (1974) observed this behaviour in the short tailed hawks of Florida. Generally, these hawks would adopt soaring, either 'balancing' in a fixed position on deflected air currents or spiralling upward in small circles. Following an unsuccessful attempt to capture a prey item, they execute an alternating flap and sail direct search across the treetops for a short distance, endeavouring to relocate the prey before resuming their initial search strategy. Interestingly, this direct search technique was amongst a selection of initial search techniques, which included sit-and-wait (perched in trees), soar-and-stoop (high level soar followed by a stoop) and soar-and-sneak (high level soar followed by an inconspicuous direct line movement), adopted by a variety of Florida based hawks and described by Johnson and Peeters (1963), but was not observed as an initial technique by Ogden.




Figure 1. Indicative smooth pursuit (SP), tracking saccades (TS) and interception saccades (IS).

Usefully, and in support of techniques for tracking visually occluded targets, it can be argued that existing ethological research into how predators visually track targets can be applied to the problem of reacquiring targets after detection has been lost. The rationale being that when animals are visually tracking a target they must apply predictive techniques to overcome intrinsic sensorimotor delays, and therefore, tracking occurs in a cycle of smooth pursuit and predictive saccades. Orban de Xivry and Lefèvre (2007) provide insight into how the two elements combine in a single sensorimotor process in humans. In this case the smooth pursuit operates with a fixed delay and maximum tracking speed, and can operate successfully in scenarios that change within these boundaries. However, where abrupt changes occur or targets are moving quickly, corrective saccades are applied, albeit at the cost of visual acuity during the movement. It is important to note that saccadic movements do not utilise visual feedback from the target, they are totally predictive and are intended to place the target back into the fovea (the region of the retina with the highest acuity). The relationship between the two modes is dependent on the ability of a smooth pursuit system; for example, Orban de Xivry and Lefèvre (2007) compare the human process with that of cats and note that the latter rely more heavily on saccades due to the lower and more variable gain in their smooth pursuit feedback system. Klam *et al.* (2001) reported two saccadic techniques in the sensorimotor process of cats: the tracking saccade and the interception saccade. The former operates in much the same way as for humans and is utilised where target speeds are lower or positional error is deemed small. Interception saccades are more extreme in their movement and intentionally overshoot the predicted target position before attempting to operate smooth pursuit. Figure 1 shows indicative movements for all three strategies. In the first third of the graph, the sensorimotor process can keep track with the target; in the middle third, the target speeds up but the process can keep track by making small catch up tracking saccades; in the final third, where target speed increases still further, interception saccades are made. The defining features of the overall process are that the least time possible is spent in saccadic motion, where visual acuity is low, and that the target is kept within the fovea for the maximum time during smooth pursuit.

Crosswind casting is a physical movement pattern employed by many insect species that use chemical odours to locate prey, for example, fruit flies (Budick and

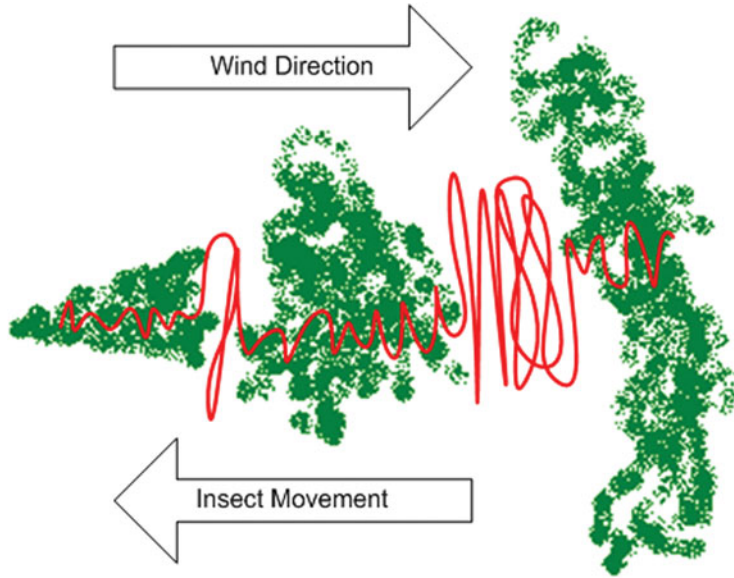


Figure 2. Indicative movement path of an insect (red trail) through an odour plume (green).

Dickinson, 2006) and moths (Willis and Arbas, 1998). During anemotactic search in odour plumes, insects employ a tacking motion whilst surging upwind toward the source. The nature of odour plumes is such that the scent is often inconsistent and, when a gap is encountered, the insect turns perpendicular to the wind direction and crosswind casting is applied to reestablish a lost odour trail (Figure 2).

2.2. *Knowing when to give up.* Once a predator has assumed a search tactic, it must then decide when to switch to another (for example switching back to the initial search behaviour from an unsuccessful attempt to reacquire a target) or, during saltatory search, move location (O'Brien *et al.*, 1989). This is known as the giving up time (GUT). The decision could be based on factors such as estimated prey numbers, capture rate, or time (Bell, 1991). For example, juvenile plaice (Hill *et al.* 2000) use a short term klinokinetic strategy; that is, they will only exploit a patch of prey for a short time before swimming away (by reducing their turn rate). This, it was argued, is because other external pressures are operating; if they loiter to fully exploit the patch, there is an increased possibility of predation.

For this work, a simple chronological mechanism is applied to return the agents to their initial search strategy following an unsuccessful search although more elaborate mechanisms may be beneficial and therefore may provide an area for further work. For example, Nishimura (1999) examined optimal GUT in uncertain environments, the findings of which could be used to develop improved GUT strategies for agents searching for faster targets.

3. **REACQUISITION STRATEGIES.** The strategies implemented in this research that deal with reacquisition search are termed Predictive Memory System (PMS) strategies, since they attempt to predict the most likely location of a lost

target, utilising a short-term memory and simple trigonometry. The PMS, regardless of strategy, has several elements:

- *Target location prediction.* Targets are either predicted to have remained in the vicinity of the last known location or to have continued on their last known heading and speed. Both scenarios model behaviours observed in nature (see Section 2.1): the former is included to assess whether it could be of value against targets that have a high turn rate (and hence prediction could be less reliable); conversely, the latter is incorporated to assess whether it could add value where targets are more predictable.
- *Reacquisition attempt decision.* To facilitate cooperation and, critically, interception, agents share information on observed targets. All agents, on noting that they are no longer receiving information regarding a previously communicated target, have to decide whether they should attempt to intercept the target based on a predicted position. Such a decision is either based on membership of a given neighbourhood or the ability to reach a predicted target location within a given time period (*i.e.* no later than when the target is predicted to pass through).
- *When to give up.* Once an agent decides to search for a lost target it should only continue to do so for a finite period of time, otherwise it may be wasting energy seeking a target that is no longer present or has deviated significantly from its predicted behaviour. Even where agents do not initially decide to engage in a reacquisition search they should still maintain the target's predicted position in memory for a period of time in case the agent finds itself in such a position that to change to a reacquisition search may be profitable. In this implementation a single giving up time is used because the likelihood of redetection diminishes at the same rate (*e.g.* an agent finding itself entering a neighbourhood after 50% of the giving up time has passed will only attempt reacquisition for the remaining 50%).

The specific PMS strategies are now considered, starting with those where the decision to attempt reacquisition is based on membership of a neighbourhood.

When a target is lost, the PMS neighbourhood determines which agents may be associated with attempts to reacquire the target. Two neighbourhood shapes were implemented: circular, for predictive tracking and direct search; and V shaped for predictive tracking and crosswind casting (hereafter, the two predictive tracking variants are termed predictive-C and predictive-V, respectively, to avoid confusion).

Neighbourhoods generally follow the underlying metaphor on which the reacquisition strategy is based. A circular pattern only was applied to direct search since that approach considers that the target is still located somewhere in the vicinity of its last known location; that being the case, any agent within a reasonable distance may be able to reacquire it. Similarly, crosswind casting applies only a V shaped neighbourhood since the underlying metaphor is based on tracking odour plumes and assumes that they continue to travel in the same direction, in which case agents located behind the target would not expect to be able to redetect since the target would continue to travel away from them. Predictive tracking – *i.e.* the smooth pursuit described by Orban de Xivry and Lefèvre (2007) – also expects the target to maintain its speed and direction; however, it was also implemented in a circular




Figure 3. Scenario where predictive-V PMS is applied.

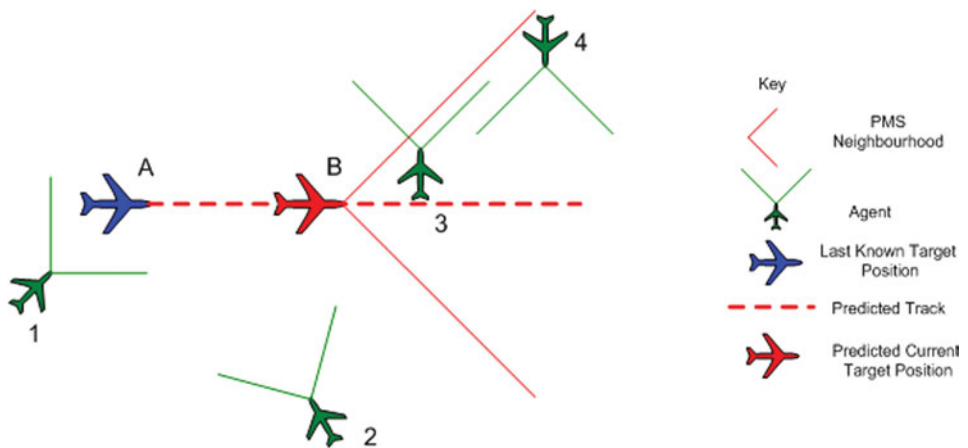


Figure 4. Scenario where crosswind casting PMS is applied.

neighbourhood to reflect an ‘optimistic’ tracker that may continue to track the target in the hope that it would change direction or slow down.

Consider the scenario depicted in Figure 3. When target detection is lost (at point A) the agents continue to predict its path. If, at any time until the GUT has expired, an agent finds itself in the neighbourhood, it attempts to re-detect the target by moving toward its predicted actual location. No lead was included, since to do so would begin to move away from the predictive search metaphor into that of a saccadic search. At point B, agent 3 finds itself in such a position and, since predictive-V PMS is in operation, the agent moves toward the predicted location.

In the next example, Figure 4, a similar situation exists except crosswind casting is in operation. By the time the target, assuming it follows its predicted path, reaches point B agent 3 has been inside the neighbourhood for some time, has crossed the predicted path and the target will pass behind it; Agent 4, however, has only just joined the neighbourhood and may well be in a better position to intercept. An obvious solution to the problem of agents passing in front of a target and not being in a suitable position to re-detect would be to allow the agents to slow down such that

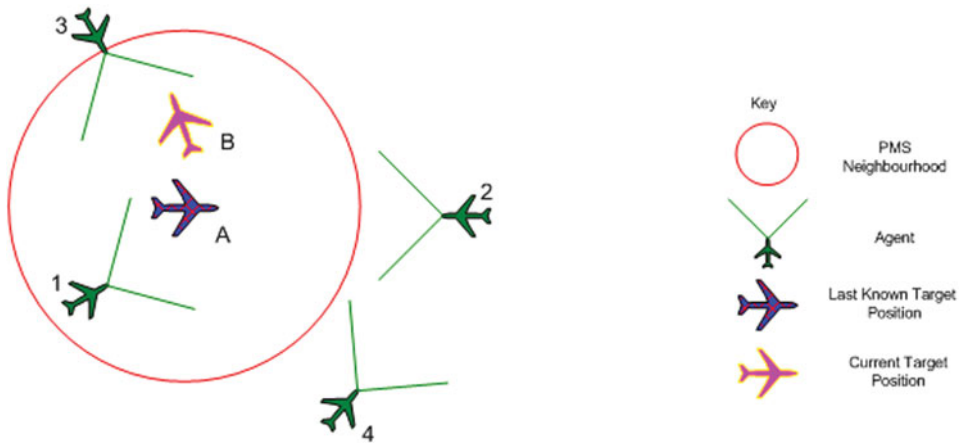


Figure 5. Scenario where direct search PMS is applied.

they do not pass the predicted path before the target. However, this was not implemented because it would move too far from the natural behaviour. Moreover, it may not be bad behaviour to operate as implemented because the target could deviate from its path and even if it does not, the agent could reacquire when it turns back toward the predicted path to continue casting.

Figure 5 depicts a direct search scenario, where the agents, finding themselves in the neighbourhood, move toward the last known position (point A). In this case the target has moved and is now at point B and, therefore, despite the unpredictable target behaviour, agent 3 may be able to reacquire. Predictive-C neighbourhoods behave in a similar manner except the locus of the neighbourhood moves with the predicted target location; agents finding themselves behind the target optimistically pursue the target in case it turns round or slows down.

This optimism, suggested by agents using the predictive-C circular neighbourhood, was not conferred on the crosswind casting strategy since the agent turns at right angles to the target movement and, therefore, if the agent were behind the target to begin with, it would become increasingly unlikely to redetect.

The dynamic behaviour of the neighbourhood could be realised through a variety of linear and non-linear functions. For this work, a simple but effective linear function was applied to control the length of the conical neighbourhood with a fixed angle, θ .

$$length(l) = r \cdot c \cdot s \tag{1}$$

Where r is the sensor range, c is a constant and s is the relative target speed. The angle, θ , and constant, c , are tuned, along with the GUT to provide good performance; Table 1 provides a summary of tuning results.

An agent may find itself in multiple neighbourhoods, in which case the selected target is resolved based on target fitness, $f(tgt)$, following Equation 2, which for the purposes of this work is based on the likely success of launching a guided weapon (as previously used in Banks *et al.*, 2008).

$$f(tgt) = \frac{1 + \cos \psi}{dist} \tag{2}$$

Table 1. Summary of tuning experiment results.

	c	θ	GUT
Circular neighbourhood	1.0	—	50 steps
V shaped neighbourhood	2.5	80°	350 steps

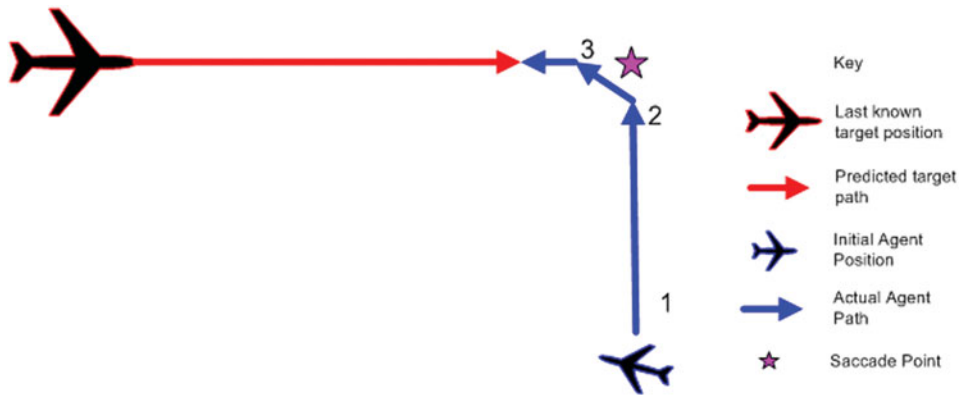


Figure 6. Example path of an agent intercepting an occluded target using a saccadic PMS.

Where ψ is the reciprocal of the target heading with respect to the agent heading and $dist$ is the Euclidean distance between the target and the agent. A minimum neighbourhood size was also included for both shapes and set such that, where a target has a relative speed of less than equality, the neighbourhood was set to that of an equal speed relationship. This was introduced because an agent losing detection of a slow target could have found itself outside the neighbourhood straight away, precluding a reacquisition search; for example, if the target was lost at the edges of sensor range then, due to a low relative target speed, the neighbourhood could be smaller than the agent's sensor range. In such cases, with the minimum neighbourhood, the agent would switch back to initial search without attempting reacquisition. For the circular neighbourhood, Equation 1 determines the radius of the neighbourhood in relation to the relative target speed and sensor range.

Saccadic PMS (Figure 6) does not apply a neighbourhood, but imitates the sensory behaviour employed where a predator is unable to continuously detect a target due to sensory limitations (the agent moves in the same way as the point of focus in the natural metaphor). When a target becomes lost (point 1) all agents predict the target path and identify a location termed the saccade point where it (the agent) could intercept the path in the shortest time, *i.e.* the closest intersection. If the agent calculates that it can reach the saccade point before the target it will continue with the attempted reacquisition, otherwise it will continue with its initial search behaviour. Should the agent reach a location within sensor range of the saccade point (point 2), and the target is not there, several scenarios could exist such as the target could have sped up, deviated from its predicted path, slowed down or simply not reached the location yet. In keeping with the natural metaphor, the saccadic agent will assume one of the latter two scenarios and will turn to track back along the predicted path




Figure 7. Indicative movement pattern produced by random walks, left: Brownian motion search; right: straight line search.




Figure 8. Indicative movement pattern produced by, left: saltatory search (note: dark points on route highlight points where the agent paused and searched intensely); right: fruit fly search.

toward the last known position. Like direct search, the PMS is ceased if either the PMS GUT is exceeded or the last known position reached.

4. ENVIRONMENT AND OCCLUSION MECHANISM. The reacquisition strategies outlined in Section 3 were implemented in a simplified model constructed in Ada 95 using the John English Windows Library (JEWL) (English, 2000) and AdaCore’s Gnat Programming System (AdaCore, 2005), as used in Banks *et al.* (2009) with indicative movement patterns reproduced in Figures 7, 8 and 9 from the movement behaviours that were also utilised in this work. To summarise:

- *Random walks (straight line and Brownian motion).* To model the stochastic leg length (l_j), which has continuous probability (P), described by Viswanathan *et al.* (1999) as:

$$P(l_j) = l_j^{-\mu} \tag{3}$$

the leg lengths were set by truncating the result of $^{-\mu}\sqrt{\text{rand}}$, where rand is a uniformly distributed random number in the range [0,1] and μ sets the path




Figure 9. Routes followed by deterministic targets to simulate an attack by the targets on the two non-central locations within the search area.

tortuousness, *i.e.* $\mu=1$ for straight line and $\mu=3$ for Brownian motion. Note: straight line search also included behaviours reported by Bell (1991) and, for straight line only, the random numbers were drawn from the range $0.0005 < \text{rand} \leq 1$ in order to constrain, for pragmatic reasons, the leg lengths, since values less than 0.0005 result in excessively large legs. (Figure 7)

- *Saltatory patterns.* Saltatory movement applied Lévy path search during the intensive search phase, *i.e.* a random walk with $\mu=2$, but also slowed the agent down by 50%. A modification to the parameters provided in the literature was the reduction in the GUT of the intensive search. This was made in acknowledgement of the environment, the targets were not hidden and to linger for too long (where there was plainly no target) would not have mimicked natural predator behaviour. (Figure 8)
- *Fruit fly.* This utilised representative movement patterns, *i.e.* overall patterns produced mimicked those presented in the literature (for example, Tammero and Dickinson 2002), since the systems and stimuli to switch between phases of flight, such as the visual feedback from the environment described in the literature (Frye *et al.*, 2003), are absent in this application. (Figure 8)
- *Deterministic movement.* Targets followed one of four routes across the search area, simulating a simultaneous attack against two non-central target locations (marked target 1 and 2 in Figure 9).

Once again, the search area was a two-dimensional representation of an area measuring 400×300 units about which the agents travelled at a constant speed of 0.8 units per step, turning at a maximum rate of 10° per step, and had a sensor range of 40 units with a Field Of View (FOV) angle of 40° . These constraints were applied to ensure fairness in comparing approaches, although in the natural world search strategies are often physiologically limited and the same would apply to physical implementations of naturally inspired systems. The comparative metric used throughout this work is the mean time to locate and destroy targets (targets had to be within 10 units and $\pm 20^\circ$ for agents to destroy them), where each time unit is a

simulation step; that is, all agents would have moved 0.8 units and all targets would have moved their respective relative distance, for example, at relative target speed of 1.5 the targets would move 1.2 units.

Since the simulation environment is featureless, it was necessary to implement an occlusion mechanism to simulate the target becoming occluded from a detecting agent, thereby facilitating evaluation of the PMS in scenarios where targets are lost whilst still within sensory range. Two indicative occlusion types were identified. The first, permanent occlusion simulates a target becoming occluded from the system due to an agent failure (such as sensor failure). When permanent occlusion is applied, it always occurs upon the first detection by an agent and the target remains occluded to the initial detecting agent but the target remains detectable by other agents. The second type, temporary occlusion, simulates where targets may become occluded through environmental features or target evasion manoeuvres. In this scenario the target occlusion is not certain and occurs randomly with 5% probability on each stepwise move. Redetection is possible by the original detecting agent, once again randomly, except with a 1% probability on each stepwise move. These figures were selected to offer a challenge to the system and therefore allow the merits of the different PMS strategies to be evaluated. As with permanent occlusion, other agents could detect the target whilst it remains occluded to the agent that first detected it.

5. EXPERIMENTS AND RESULTS. The five PMS strategies were applied to 10 agents that were pitted against four dynamic targets having a range of relative speeds {0.5, 1.5} and movement patterns {Brownian motion, deterministic, fruit fly} to examine the effect of maintaining target memory on overall performance against occluded targets. Upon initial detection, the targets were occluded from the detecting agent either permanently (permanent occlusion), simulating agent sensor failure during detection, or temporarily (temporary occlusion), which simulates temporary obscuration of the target to the detecting agent.

5.1. Permanently occluded targets. Against Brownian motion targets (Figure 10) the strategies that depended less on the target being in the predicted position (direct, predictive-C) generally performed better. It can also be seen that the V shaped neighbourhood of predictive-V was stronger when used with more ranging types of initial search. This variation was caused by the agents being moved toward the target during PMS and then, following GUT expiry, continuing to move quickly into nearby areas where the lost target had moved to. This same rationale explains the direct and predictive-C strategies' competitive performance across the range of relative speeds; at lower speeds they ensured that cooperating agents were directed to the general area from a range of directions (although moving them away from areas that may have contained as yet undetected targets), whilst, conversely, at higher speeds the prediction of the occluded target was more likely to be inaccurate and thus agents would have been directed away from the actual target location but the extended neighbourhood area increased the likelihood of the target being re-detected.

Where target behaviour was more predictable, as for deterministic targets (Figure 11), the strategies that only committed agents that were likely to be able to re-detect held the advantage, whilst techniques that committed more agents that were




Figure 10. Summary of overall target reacquisition performance against four Brownian motion targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where permanent occlusion occurs.

unlikely to succeed, *i.e.* the circular neighbourhoods (predictive-C and direct search), were less successful because they needlessly moved agents away from areas that may have contained other, undetected, targets.

The predictive-V strategy did not perform well against fruit fly targets (Figure 12) because, despite the target movement type's long inter-saccade leg lengths, the unpredictable heading changes and switches to the short saccades mean the target is unlikely to be in the vicinity of its predicted location. This was in contrast to the other predictive neighbourhood, predictive-C, which, although not the best, remained competitive due to the target's propensity to switch into its intensive search mode (and therefore agents being drawn in from behind the target's last known position would be likely to detect it on their way to the predicted position, as was the case against Brownian motion targets). Agents applying saccadic reacquisition search were successful for similar reasons, except where targets remained on a long leg the agents that were able to intercept them changed their route to do so; where the target switched to intensive search – or was already in that phase when they became occluded – the agents would have moved back along the predicted track to the intensive search area, thus increasing the likelihood of re-detecting the target.

5.2. Temporarily occluded targets. When occlusion was temporary, the direct search became more prominent against Brownian motion targets (Figure 13), because the agent that last detected the target maintained its focus on the area where the target was lost; in Brownian motion targets this was likely to remain a




Figure 11. Summary of target reacquisition performance against four deterministic targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where permanent occlusion occurs.

good area to search because the target would not move far away from the area. The predictive-V also became more prominent since it committed fewer agents to the search.

Against deterministic targets (Figure 14), the results remained similar to permanent occlusion due to the predictability of the targets, although the direct search strategy did not perform as strongly for the same reason that it improved against the Brownian targets: the losing agent remained in the area, this time unproductively, when it could have moved on to detect other targets.

Finally, where the targets once again present an element of intensive searching, as found in fruit fly movement, the memory systems that lead agents to the memorised location, whether that be the last known or predicted new location, perform best (Figure 15).

6. DISCUSSION. During permanent occlusion, and where the target was following Brownian motion, it was seen that all PMS types could improve performance over using no memory. Intuitively, for Brownian motion targets one might have expected a direct strategy to be the best performer since the target would not move far from the last detection location. However, the agent to target ratio was low and therefore the neighbourhood shape could commit a relatively large number of the agents to the reacquisition search. In the predictive-V, although the target may




Figure 12. Summary of target reacquisition performance against four fruit fly movement targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where permanent occlusion occurs.

have changed direction it would not have moved far, therefore the co-operating agents that had ranging initial search behaviours would move into the right area and as such, even where the reacquisition search was unsuccessful, the switch back to the initial search could re-detect the target. The crosswind and saccadic systems suffered because of the unpredictable nature of the target; the former would not move the agents toward the target's location and the latter altered the agents' route down the route least probable to encounter the target for the target movements adopted (although, again, where used in conjunction with a ranging search this was not so problematic).

Against deterministic targets, the predictive-V dominated since the agent usually predicted the target's path correctly. Crosswind and saccadic systems performed relatively better because the target was now moving toward the agents. Unfortunately, for the crosswind system the agents were often not facing the target as it passed by. In nature, crosswind casting is used to re-detect targets that have a broad signature (*e.g.* a plume of odour); by changing to a perpendicular movement the animal increases its chance of regaining detection. Where the target is relatively small, as in the simulation, such behaviour actually reduces the possibility as the target passes by. A further noteworthy point to be made regarding the positioning of the agent during crosswind casting (in a real implementation) is that it could place the vehicle in a position vulnerable to counter-attack by the target since the search places the vehicle perpendicular to the target track. The direct and predictive-C




Figure 13. Summary of overall target reacquisition performance against four Brownian motion targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where temporary occlusion occurs.

systems once again suffered due to the over commitment of system resources, with the latter performing better due to the generally correct prediction of the target location.

The surprisingly poor performance of the V shaped system against fruit fly targets where permanent occlusion occurred could be regarded as a warning about over-tuning a system to a particular type of target. During the tuning, the GUT was found to be very short against fruit fly targets but very long against deterministic targets. Since the agents did not have a learning mechanism that allowed them to select an appropriate GUT based on the target movement type they applied an inappropriately long amount against targets that were often not following the path they were last detected on. This was not so problematic for the saccadic system since those agents gave up after finding that the target was not following the predicted path from its last known location. However often the fruit fly targets did apply long leg lengths which meant that the saccadic system was worthwhile. The other three systems, whilst better than nothing, still suffered from the same issues as discussed for other target types (over commitment, agents facing the wrong direction, *etc.*).

When an agent could re-detect the target it had lost, results changed significantly. Against Brownian motion targets, the direct approach worked best because the agent remained searching in the area that it last detected the target and when the occlusion




Figure 14. Summary of target reacquisition performance against four deterministic targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where temporary occlusion occurs.

ceased it was able to attack; of course, the same was true for permanent occlusion except that re-detection never occurred and the opportunity was lost. Where the PMS predicted target location, performance was only improved significantly either where agents were drawn in from all directions (predictive-C) or where, for predictive-V, the initial strategy could move the agents quickly away from the incorrectly predicted location to the actual location nearby (more so at slower target speeds, since the target would not have moved so far away).

Continuing with temporary occlusion scenarios, where deterministic targets were operating, the predictive-C system held the advantage over the predictive-V system because when an agent detected a target that was moving away from it, remembering that results include targets moving more slowly than agents, it remained in the reacquisition neighbourhood; in the predictive-V system that was not the case, and even when an agent was capable of 'chasing down' a target it had lost, it would not attempt to because it was not in the qualifying neighbourhood. The remainder of the systems held their relative positions.

In the final occlusion based PMS experiments, against fruit fly targets, where temporary occlusion occurred, the circular neighbourhoods were best performers. Again, this was because of the inter-saccade phase where targets did not range far (in the case of direct search) and the ability of the predictive-C scheme to chase down targets travelling away from them at the time of occlusion. V shaped neighbourhood systems did not suffer as much where targets could be reacquired by the losing agent,




Figure 15. Summary of target reacquisition performance against four fruit fly movement targets moving at 0.5 times (top) and 1.5 times (bottom) agent speed where temporary occlusion occurs.

because quite often the agent would be in the neighbourhood and be able to reacquire the target before it moved out of sensor range.

7. CONCLUSIONS. This work set out to extend the authors' previous study on nature inspired search behaviours to incorporate additional behaviours to tackle instances of occlusion – whether temporary (e.g. due to physical obstruction) or permanent (e.g. due to sensory failure). The application of nature inspired predictive memory systems was shown to improve systemic capability, although the agent–target relationship was seen to be important in how much improvement was gained. Where targets behaved predictably, then predictive tracking only applied by agents finding themselves in a V shaped neighbourhood performed best. Less predictable targets that moved away from the area quickly were best dealt with by predictive tracking with circular neighbourhoods or saccadic behaviour. However, the application of a PMS was seen to be inefficient when applied unnecessarily, since, once committed to the pursuit of a memorised target, the likelihood of an agent detecting other targets is diminished due to the reduction in the turn rate of the agent; i.e. the agent starts to move in a straight line to intercept the target, which, as found in our previous work (Banks and Vincent, 2009), is generally less effective in scenarios with moving targets. Hence, the importance of understanding the target–agent relationship present in a given scenario, or implementing some form of adaptive mechanism to learn from experience of target behaviour.

REFERENCES

- AdaCore, (2005). Ada academic initiative. Available from: http://www.adacore.com/academic_overview.php [accessed 26 January 2005].
- Banks, A., Vincent, J. and Phalp, K., (2008). Particle swarm guidance system for autonomous, unmanned aerial vehicles in an air defence role. *The Journal of Navigation*, **61**, 9–29.
- Banks, A. and Vincent, J., (2009). Exploring the performance of natural search strategies for the control of unmanned autonomous vehicles. *The Journal of Navigation*, **62**, 283–301.
- Bell, W. J., (1991). Searching behaviour: the behavioural ecology of finding resources. *Chapman and Hall Animal Behaviour Series*, Chapman and Hall, London.
- Budick, S. A. and Dickinson, M. H., (2006). Free-flight responses of *Drosophila melanogaster* to attractive odors. *The Journal of Experimental Biology*, **209**, 3001–3017.
- English, J., (2000). About JEWEL. Available from: <http://www.it.bton.ac.uk/staff/je/jewel/> [Accessed 22 June 2005].
- Frye, M. A., Tarsitano, M. and Dickinson, M. H., (2003). Odor localization requires visual feedback during free flight in *Drosophila melanogaster*. *The Journal of Experimental Biology*, **206**, 843–855.
- Hill, S., Burrows, M. T. and Hughes, R. N., (2000). Increased turning per unit distance as an area-restricted search mechanism in a pause-travel predator, juvenile plaice, foraging for buried bivalves. *Journal of Fish Biology*, **56**, 1497–1508.
- Johnson, N. K. and Peeters, H. J., (1963). The systematic position of certain hawks in the genus *Buteo*. *The Auk*, **80**, 4, 417–446.
- Klam, F., Petit, J., Grantyn, A., Berthoz, A., (2001). Predictive elements in ocular interception and tracking of moving targets by untrained cats. *Experimental Brain Research*, **139**, 233–247.
- Nishimura, K., (1999). Exploration of optimal giving-up time in uncertain environment: a sit-and-wait forager. *Journal of Theoretical Biology*, **199**, 321–327.
- O'Brien, W. J., Evans, B. I. and Browman, H. I., (1989). Flexible search tactics and efficient foraging in saltatory searching animals. *Oecologia*, **80**, 100–110.
- Ogden, J. C., (1974). The short-tailed hawk in Florida I. Migration, habitat, hunting techniques, and food habits. *The Auk*, **91**, 95–110.
- Orban de Xivry, J. J. and Lefèvre, P. L., 2007. Saccades and pursuit: two outcomes of a single sensorimotor process. *Journal of Physiology*, **584**, 11–23.
- Tammero, L. F. and Dickinson, M. H., 2002. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *The Journal of Experimental Biology*, Vol. **205**, pp. 327–343.
- Viswanathan, G. M., Buldyrev, S. V., Havlin, S., Da Luz, M. G. E., Raposo, E. P. and Stanley, H. E., 1999. Optimizing the success of random searches. *Nature* **401**, pp. 911–914.
- Willis, M. A. and Arbas, E. A., 1998. Variability in odor-modulated flight by moths. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, vol. **182**, pp. 191–202.