3D SIMULATION OF A SENSORIMOTOR STEALTH STRATEGY FOR CAMOUFLAGING MOTION

Andrew James Anderson and Peter William McOwan
Department of Computer Science, Queen Mary, University of London, E1 4NS, UK
aja@dcs.qmul.ac.uk

ABSTRACT

A sensorimotor control system for a biologically inspired stealth strategy (motion camouflage) intended to conceal the motion of a predator from its prey is simulated. The control system is formed from three Multilayer Perceptrons trained using Backpropagation. In the simulation the control system, operating using realistic input information, is shown to be able to track prey moving in 3D space. This extends previous work that has only considered two dimensions.

1. INTRODUCTION

Motion camouflage is a stealth strategy that allows one moving body to camouflage its motion from another moving body. The technique was first suggested in [1], where it was observed that, in mating, male hoverflies may employ motion camouflage to sneak up on females. In [2] artificial neural control systems for motion camouflage were simulated operating in two-dimensional environments. They were shown to be able to accurately approach a prey that moved along either artificially generated flight paths or flight paths recorded from real hoverflies. Also it was shown, in both cases that it was possible to improve performance through prediction. This paper covers the progression of the simulation to three dimensions and the subsequent testing of performance.

2. MOTION CAMOUFLAGE BACKGROUND

Throughout the following discussion the predatory body that wishes to remain camouflaged shall be termed the *shadower*, and the body being tracked, the *prey*. The initial position of the shadower in the pursuit shall be referred to as the *fixed point*. The basis of motion camouflage is that the shadower should approach the prey in such a way that it appears to be a stationary object in the environment. This is achieved by the shadower ensuring that it is always positioned directly in between the fixed point and the prey. For instance if the shadower were to start its approach positioned in front of a rock, it would ensure that it is always positioned directly in between the rock and the current position of the prey. The optic flow of the shadower projected onto the retina of the prey would then emulate that of the rock. In other words, the prey would always see the shadower silhouetted against the rock, and it is hoped not be alerted to the approach of the shadower. Motion camouflage trajectories generated by the control systems trained in [2] tracking a hoverfly trajectory and an artificially created regular trajectory are displayed in **Fig.1**.

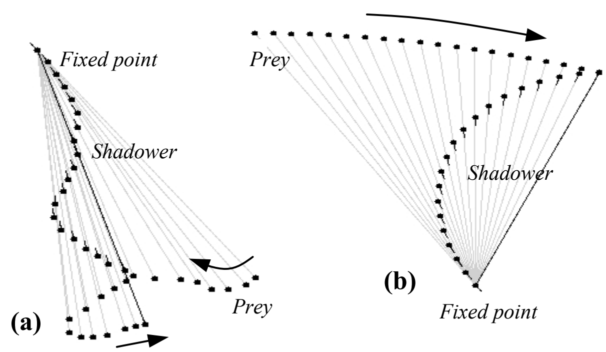

Fig. 1. Camouflage trajectories generated by the control systems of [2], tracking **(a)** the flight path of a real hoverfly **(b)** the prey moving along an arc. The shadower is depicted by the dots with tails, the prey, the dots without. At each instant the shadower is expected to lie on the camouflage constraint line joining fixed point and prey. The final constraint line is highlighted in bold

In order to be able to calculate camouflaged movements the shadower must have a concept of its position relative to its initial position and the prey. In 2 dimensions it was shown that the following input information was sufficient to make this computation:

- The current direction of the prey;
- The image motion of the prey (i.e. the change in direction of the prey between time steps);
- A memory of recent control system output (movements).

The movement trained was the direction in which to move and rotation to turn. The shadower was trained to rotate to point radially away from its starting position as suggested and observed in the hoverfly behaviour of [1]. This rotation implicitly acts as an estimation of the direction of the fixed point. If the shadower rotates accurately it knows that the fixed point is directly behind it (whilst the prey should be directly in front).

In [2] it was seen that in order to perfectly camouflage its approach, the shadower would have to make some sort of prediction of future prey movement. For instance, if the shadower moves so as to be camouflaged according to the time instant at which it commenced its movement, on completion of the movement the prey may have moved elsewhere. Thus, unless the shadower is able to respond quickly enough, its camouflage is likely to be broken. The first method of approach was termed the *predictive* algorithm and the second the *responsive* algorithm.

It was shown in [2] that the performance of the predictive artificial neural control systems was closer to perfect camouflage than that produced by the responsive algorithm (calculated using trigonometry from the exact positions of shadower, prey and fixed point).

This paper covers the extension of the previous 2D control systems to operate in a virtual 3D environment. Similar tests are carried out upon the controllers in order to ascertain whether they benefit from any predictive capability. Beyond this it is investigated whether the control systems are able to operate without knowledge of prey image motion. Results are presented in **Section 7**.

3. 3D CONTROL SYSTEM DESIGN

The design of the control systems follows that of [2] (see **Fig. 2**). Each controller comprises three Multilayer Perceptrons. The first of these networks is responsible for estimating the distance to the fixed point. The second, the direction to move and the third, the rotation.

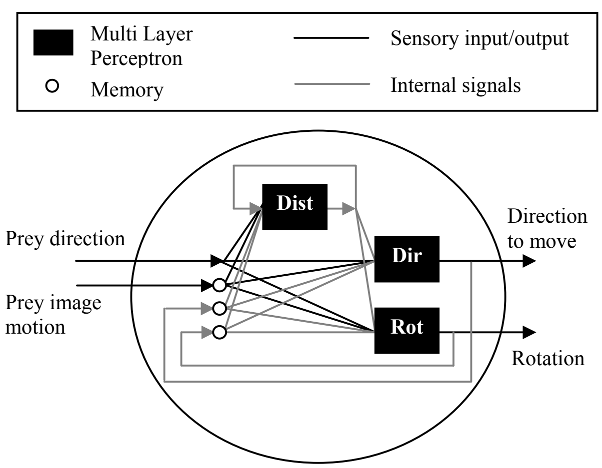

Fig. 2. Control system architecture. See text for explanation.

At each time step the control system is supplied with the following sensory input (Note that as described in **Section 6**, not all of the control systems were given image motion as input):

- The direction of the prey, represented as two angles (in a spherical coordinate system relative to the shadower as origin) see **Fig. 3**;
- The image motion of the prey over the previous 3 time steps, represented as the difference in the polar bearing and difference in azimuth between consecutive time steps.

And the following internally generated signals:

- Feedback of the previous direction and rotation output, both represented using two angles as for the prey direction. The previous three outputs are memorised.
- An estimate of the distance from the fixed point. At each time step the distance is estimated prior to the movement, and the new distance estimate fed directly to the direction and rotation networks, as well as back to the distance network, for the next distance estimate (the initial input to the distance network is 0).

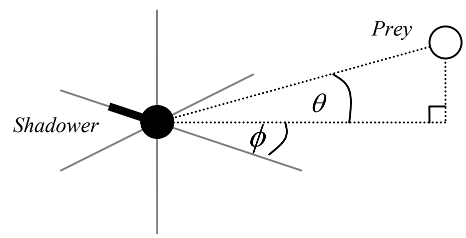

Fig. 3. Representation of prey direction, (and shadower movement direction), using a polar angle ϕ and azimuth θ.

Each of the networks was formed from 2 hidden layers, the first consisting of 40 nodes, the second 25. Each node used the logistic activation function. Training was undertaken on line using the standard Backpropagation algorithm, with a momentum term of 0.1 and a learning rate of 0.05. All parameters were selected empirically. There was not observed to be any noticeable improvement in performance or training time, through use of alternative activation functions (e.g. hyperbolic tangent, arc tangent, natural log).

4. TRAINING AND TESTING

Training and testing were accomplished by running trials in which the shadower would attempt a motion camouflage approach. In all cases the prey trajectory was predetermined. Initial shadower positions were selected at random positions within a given distance range from the prey (between 200 and 400 units in the simulation). At each step along the trial each network in the control system would be trained with the appropriate output. Target outputs were calculated *a posteriori* with knowledge of the exact locations of shadower and prey.

In the instance that the prey was in between shadower and fixed point, the shadower was trained to move directly towards the prey.

Two different measures of error were used to assess the performance of the control systems:

- *Visual error*, the angle subtended at the prey between shadower and fixed point;
- *Direction error*, the difference between the direction targets and the actual output.

Measurements of error were recorded at every step along the pursuit.

5. PREY TRAJECTORIES

Control systems were trained and tested on two types of prey trajectory, both artificially generated:

Regular trajectories, where the prey changed direction by a constant angle at each step. The change in polar angle and azimuth were independent and randomly determined before the trial. Step size was held constant at 4.5 units as in [2].

Stochastic trajectories, where the change in direction was controlled by randomly selected deviates from a normal distribution. Step size was dependant upon the current direction change. See equations (1) and (2) below, where $\Delta\phi_t$ and $\Delta\theta_t$ are the polar angle and azimuth change respectively at time t (spherical co-ordinate system with prey as origin) and r_t is the step size.

$$\Delta\phi_t = \frac{\Delta\phi_{t-1}}{p} + m \cdot z_t \qquad \Delta\theta_t = \frac{\Delta\theta_{t-1}}{p} + m \cdot z_t \qquad (1)$$

z_t is a random deviate picked from a normal distribution with mean 0 and variance 1. z_t is selected independently for $\Delta\phi_t$ and $\Delta\theta_t$. (i.e. the random component is not the same for polar angle and azimuth). p and m are terms included to modify the influence of the previous direction, and randomly chosen component on the new direction. In the simulation $p = 2$ and $m = 1/5$.

$$r_t = 2 + 5 \cdot \left(\frac{m}{m + |\Delta\phi_t| + |\Delta\theta_t|} \right) \qquad (2)$$

Thus r_t is limited to lie within the range 2 to 7 units, with a mean of approximately 4 units.

All prey trajectories were limited to 40 steps and in both cases the step size of the shadower was held constant at 5 units (just greater than that of the prey). These parameters replicate those of [2].

6. EXPERIMENTAL PROCEDURE

Beyond the successful training of a 3D control system, the experimental procedure was designed to investigate the following (interrelated) areas:

- Whether provision of image motion (see **Section 5**) as input offers any advantages;
- Whether the control systems were able to make successful predictions;
- Whether the control systems were able to generalise to the alternative prey trajectory type to those used in training.

In total, 12 control systems were trained. 6 were trained on regular prey trajectories and 6 on stochastic trajectories. In order to investigate the value of prey image motion to the calculation, half of the control systems in each group were trained with image motion as input and half without. From each group a single control system was selected for further testing. The training process and selection procedure are described below.

The controllers were trained according to the procedure described in **Section 4**. The control systems were trained for a fixed number of iterations ($5*10^8$). Throughout training their performance was tested, and a record kept of the best performing control system over the training period. The tests consisted of 100 trials, 10 trials run against each of 10 different prey trajectories. Testing parameters (e.g. initial distances and bearings of shadower from prey) were selected randomly before each test. Performance was measured using the mean directional error in polar angle. If the current state of the control system gave a lesser error than any of the previous, a second test was undertaken. If the error was still smaller, the record of the previous best controller was replaced with the current. On completion of training this record was selected to be the final state of the control system.

In order to determine which control system from each group to retain, more rigorous testing was undertaken. Each control system was tested on 10,000 trials, 100 trials run against 100 prey trajectories (of similar type to those used in training). From each group, the control system yielding the least error was then tested on the alternative prey trajectory type to that on which it had been trained. To allow testing for evidence of prediction, corresponding results for both prey trajectory types were generated using the responsive algorithm (calculated from exact positions using trigonometry, see **Section 2**).

7. RESULTS AND DISCUSSION

Camouflaged trajectories generated by the control system trained on stochastic prey movements are shown in **Fig 4**. As can be seen, the shadowers approach accurately from the three different starting positions.

The average errors resulting from the tests described in **Section 6** are displayed in **Fig. 5**. The results of statistical analyses performed on the error data are displayed in **Table. 1** and summarised below.

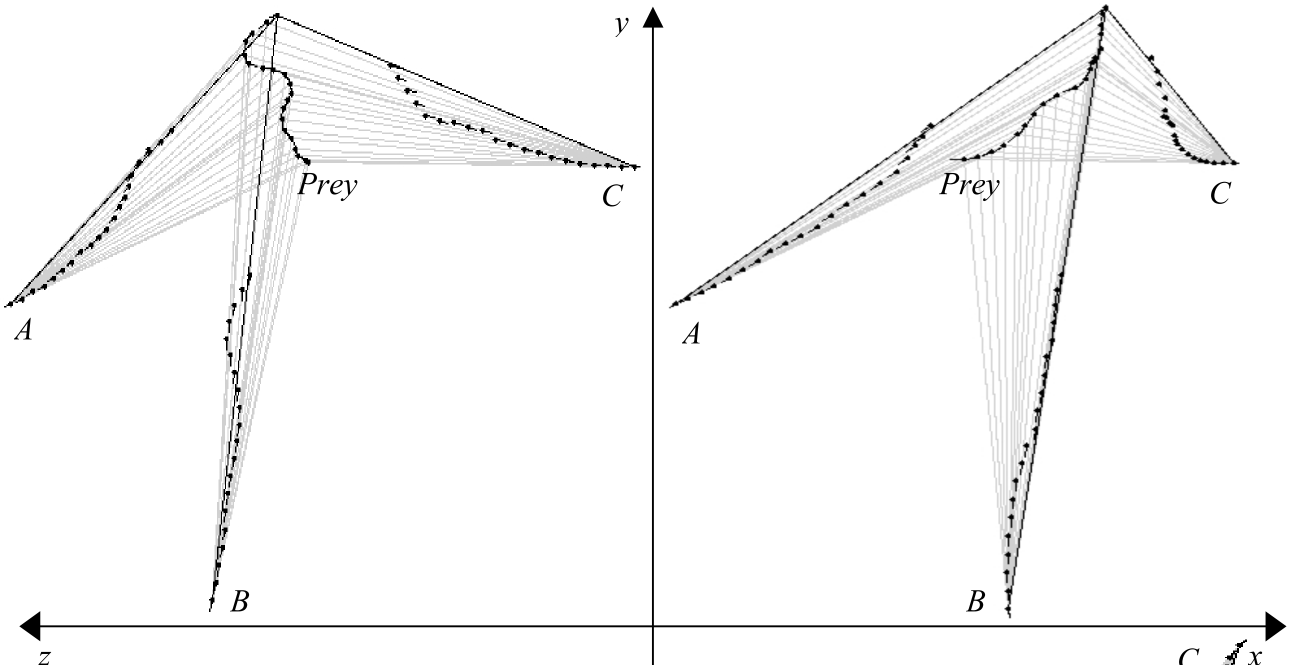

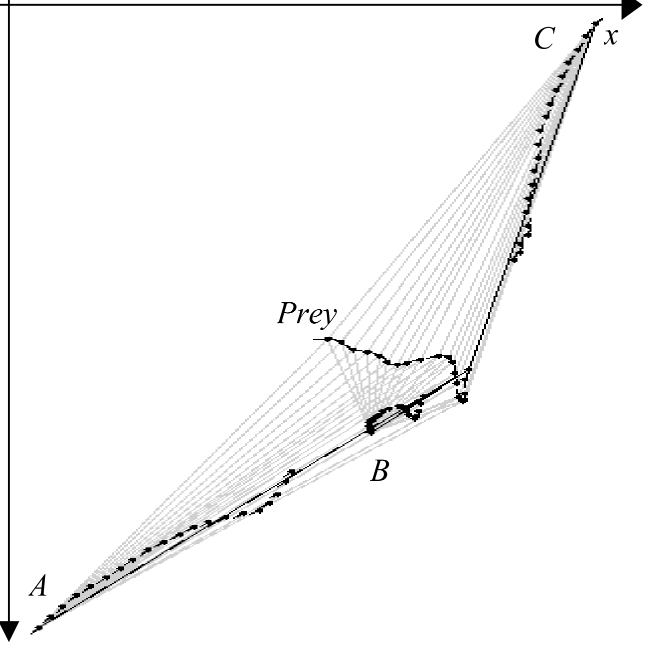

Fig. 4. Three shadowers simultaneously employing motion camouflage to approach a single prey from fixed points at *A*, *B* and *C*. Each diagram depicts the same pursuit, projected into the *xy*, *xz* and *yz* planes. Prey movement is determined stochastically as described in equations (1) and (2). As for **Fig. 1**, the final constraint line is highlighted in bold. Shadower and prey are plotted at every second time step. The control system trained on stochastic prey trajectories, without image motion inputs was used to generate the shadower trajectories.

Table. 1a displays the results of an analysis of the system trained on stochastic trajectories. It identifies a highly significant difference in visual error between tests on regular and stochastic prey trajectories. This difference can be seen in **Fig. 5**, where the error is always lesser on tests of the same prey trajectory type to those trained. There was also observed to be a significant difference in polar and azimuth error (but not visual error) between systems operating with and without image motion. The significant polar interaction results from the comparatively low error of the image motion system tested on stochastic trajectories contrasting with the slightly higher comparative error on regular prey trajectories.

All tests performed on the system trained on regular prey trajectories (**Table. 1b**) indicated highly significant differences. Conversely, this time it was the image motion system that yielded comparatively greater errors when tested on the trained trajectory type (regular) and far lesser on the alternative (stochastic).

The interaction in **Table. 1c** and **d** again highlights the superior performance of the controllers when tested on the trained prey trajectory type. The insignificant difference in visual error between training trajectory types in **Table. 1d** suggest that the image motion networks were better able to generalise to the alternative prey trajectory type (stochastic) to that trained. As such, no evidence was found to say that knowledge of image motion increases the maximum attainable accuracy of the control systems, only that it may enhance the system's ability to generalise to new prey movement patterns (when trained on regular movement).

Finally, **Table 1e** compares the error of the two non-image motion systems to that of the responsive algorithm on tests against both trajectory types (the control systems were tested on the same prey trajectory types to those trained). Again the test indicated that all differences were highly significant. In accordance **Fig. 5** shows that the average errors from the relevant controllers were lower than those of the responsive algorithm. The significant

Tested on (Regular or Stochastic prey trajectories)

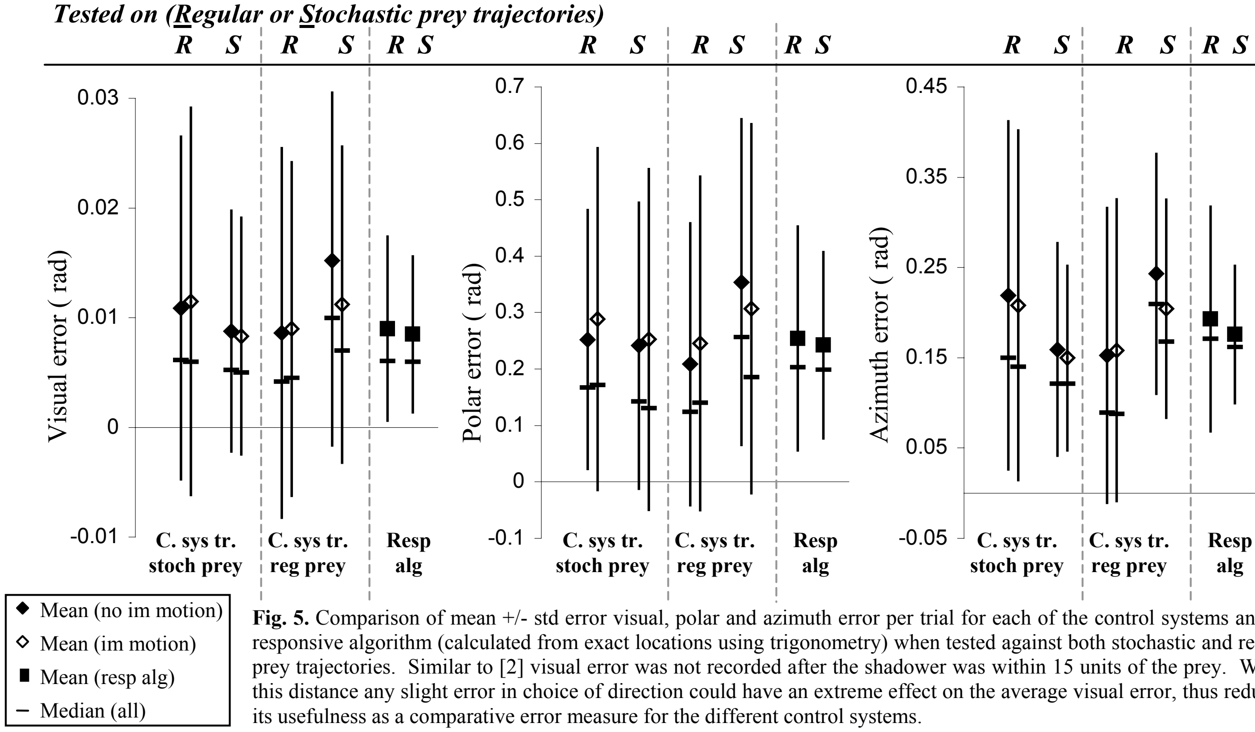

Fig. 5. Comparison of mean +/- std error visual, polar and azimuth error per trial for each of the control systems and the responsive algorithm (calculated from exact locations using trigonometry) when tested against both stochastic and regular prey trajectories. Similar to [2] visual error was not recorded after the shadower was within 15 units of the prey. Within this distance any slight error in choice of direction could have an extreme effect on the average visual error, thus reducing its usefulness as a comparative error measure for the different control systems.

Table 1a-e. Scheirer-Ray-Hare tests (non-parametric equivalent of two-way ANOVA based on ranked variates), testing for differences in error resulting from the tests described in **Section 6**. **e** investigates the difference between the responsive algorithm and control system output.

In each analysis, **bold** font indicates the parameter held constant and main effects are in *italics*. Replicates were the mean error per trial. Highly significant probabilities (P<0.0001) are shaded in dark grey. P<0.002 is shaded in light grey and insignificant probability values (P>0.05) are displayed.

		Visual error	**Polar error**	**Azimuth error**
a	**Contr sys trained on stoch prey traj**			
	Stochastic vs Regular test trajectories			
	Image motion vs without	0.4		
	Interaction	0.4		0.1
b	**Contr sys trained on reg prey traj**			
	Stochastic vs Regular test trajectories			
	Image motion vs without			
	Interaction			
c	**Trained without image motion**			
	Stochastic vs Regular test trajectories			
	Stoch vs Reg training trajectories			
	Interaction			
d	**Trained with image motion**			
	Stochastic vs Regular test trajectories			
	Stoch vs Reg training trajectories	0.1		
	Interaction			
e	**Trained without image motion**			
	Stochastic vs Regular test trajectories			
	Responsive alg vs control system			
	Interaction			

interaction is a result of the comparatively lower error of the responsive algorithm on the stochastic prey trajectory test to the comparatively higher error given by the control systems. These results show that the performance of the control systems was on average closer to perfect camouflage than the responsive algorithm. However the corresponding higher standard error (see **Fig. 5**) suggests a more erratic performance. (i.e. when mistakes were made, they were greater than those of the responsive algorithm).

8. CONCLUSIONS

This paper has simulated and tested 3D motion camouflage control systems. It has been shown that it is possible to simulate a 3D predictive control system using realistic, minimal inputs. Also that knowledge of prey image motion (as had been used in previous work [2]) is not essential for this, though, in certain situations it may be beneficial in helping the system generalise to previously unseen prey movement patterns. Future work shall investigate the application of the system to the control of autonomous agents in first person perspective computer games.

9. REFERENCES

[1] M. V. Srinivasan and M. Davey, "Strategies for active camouflage of motion," *Proc. R. Soc. Lond. B,* Vol. 259, pp. 19-25, 1995.

[2] A. J. Anderson and P. W. McOwan, "Towards an autonomous motion camouflage control system", *International Joint Conference on Neural Networks*, pp. 2006-2011, 2002.