Towards an autonomous motion camouflage control system

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

Abstract - A sensorimotor controller for a biologically inspired stealth strategy (motion camouflage) is implemented in a software simulation using back-propagation. When operating with realistic inputs, the controller allows a predator to track a prey that moves along real hoverfly flight paths, whilst appearing to remain stationary.

I. INTRODUCTION

A. Motion camouflage

Motion camouflage [1] is a stealth strategy intended to allow one moving body (a *shadower*) to conceal its motion from another moving body (the *prey*). In order to achieve this the shadower adopts a trajectory such that its image projected onto the prey's retina emulates that of a distant stationary object (the *fixed point*). This trajectory requires the shadower, at all times, to remain directly in between the fixed point and the prey, i.e. on the line (camouflage *constraint line*) connecting the prey and fixed point (**Fig. 1**). The intention of this is that as the prey always sees the shadower silhouetted against the fixed point, it will not realise that the shadower has moved. Motion camouflage could be employed to approach the prey, retreat, or to disguise movement to a particular destination. There is some evidence that motion camouflage exists in nature; male hoverflies on occasion appear to make use of the technique to track females [1]. This work concentrates on approaching the prey on a 2-dimensional landscape.

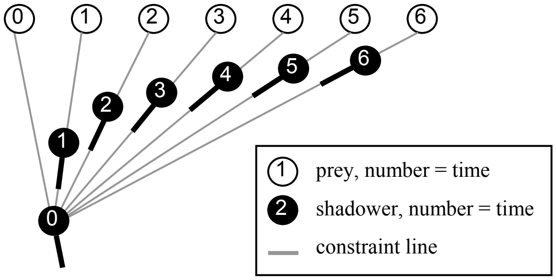

Fig. 1. A possible shadower trajectory demonstrating motion camouflage. The fixed point is located at the initial position of the shadower.

Motion will only be totally hidden when the shadower is at such a distance that the prey cannot detect any difference in its apparent size as the shadower moves. This need not imply that motion camouflage becomes useless beyond the time at which the prey notices that the fixed point has changed form. The prey still perceives no lateral motion which would be present in the vast majority of approaches. However, beyond a certain point shadower looming will unavoidably break camouflage. As such, motion camouflage is a technique designed to allow the shadower to get within a certain proximity of the prey undetected. **Fig. 2** gives an impression of the effects of shadower looming for different sizes of shadower. As the shadower nears the prey the rate of change of its apparent size increases exponentially. However at distance the gradient is minimal.

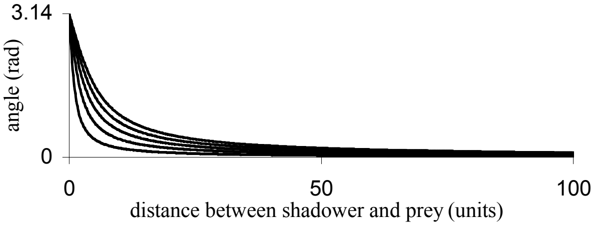

Fig. 2. The change in apparent size of shadower on approach. The apparent size is measured as the angle subtended at the prey's retina by the two outmost points of the shadower. Lines are plotted for 5 shadower widths of {2,4,6,8,10} units, with the lowest line representing a width of 2 and the upper line a width of 10.

B. Algorithms for motion camouflage

Calculating camouflaged movements is a problem of sensorimotor integration. Sensory input must be transformed into a mechanical response that provides the appropriate motion. Two conceptual algorithms are proposed here.

The first algorithm shall be referred to as the *responsive* algorithm. Adopting this, the shadower would always move so as to be camouflaged according to where it last saw the prey. But, while the shadower decides where to move and subsequently moves, the prey is liable to move away. It follows that this strategy may never involve the shadower lying on the current constraint line. However, if the shadower is able to detect very small changes in lateral prey movement over sufficiently short time intervals and respond by making corrective movements so rapidly that they are unnoticed, it will remain camouflaged.

In order to calculate each movement using trigonometry, the shadower would need to know a minimum of the following information (**Fig. 3**):

- the angle subtended at the fixed point by shadower and prey ($\Delta\theta$);
- the distance separating shadower and fixed point (ρ).
- the set of positions the shadower can expect to reach with its next movement (in **Fig. 3.** A maximum distance of $\Delta\lambda$).

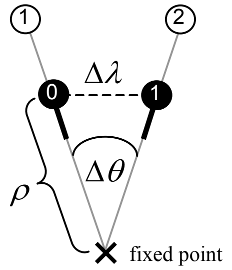

Fig. 3. Information necessary to implement the responsive algorithm using trigonometry. See text for further explanation and **Fig.1** for key.

The first problem faced by the shadower is accessing this information. Effectively the shadower must be able to estimate its position relative to prey and fixed point (i.e. the distance and direction of both). A task made more difficult if the shadower can not see the fixed point. Note that the fixed point need not be an existing landmark and could be the initial position of the shadower, in which case the prey would regard the shadower itself as a stationary landmark.

Assuming that the shadower can access and process the requisite information, the responsive algorithm gives the shadower the potential to track the prey regardless of the prey's movement pattern. However, if the shadower can not react quickly enough, the prey may notice the lag in response.

If the shadower were able to reliably predict the position of future constraint lines (i.e. implicitly predicting the prey's forthcoming movement) and move to these instead, then it should be able to improve its camouflage. This gives the second algorithm, which shall be referred to as the *predictive* algorithm. To allow prediction, in addition to the requirements of the responsive algorithm, the shadower must possess some memory of the previous positions of the constraint line/prey. To correctly retrieve this information it must also have knowledge of its own recent movements.

C. Overview

This paper presents the first virtual control system for motion camouflage known to the authors that implements the predictive algorithm. The system operates with similar levels of sensory information to that which could be expected to be retrieved by either robot or insect. This work therefore is a precursor to implementation of the system in an autonomous robot.

II. METHODS

A. Control system inputs and outputs

Emphasis has been placed on simulating sensory inputs that could be acquired by a real visual system. Although it may be possible to gain spatial information in other ways, such as hearing (e.g. owls, and actively with echolocation, bats) and electromagnetic senses (e.g. dogfish), this work concentrates on vision. Light tends to be the most accurate source of spatial information available to an animal and can be measured passively (unlike echolocation where the emitted signal could expose the shadower).

For motion camouflage to be successful, the prey must not be able to sense any difference in shadower range as the shadower approaches. This implies that the shadower must be at such a distance that the prey can not make use of binocular depth cues (e.g. binocular stereopsis, convergence, accommodation, see [2]) to judge the distance of the shadower. It follows that the shadower similarly will not be able to use these techniques to estimate prey distance. Nevertheless accurate measures of absolute depth can be gained from monocular vision. One such cue is image motion.

Praying mantids are able to estimate the distance to stationary targets from the target image motion generated as they move their heads from side to side [3]. In the case of motion camouflage the distance estimation is made more difficult because the prey also is moving. Theoretically any prey movement within a range of π radians could generate the same image motion. So, this information alone is not enough to pinpoint the distance to the prey. In practice other information is available. For example, in a given time interval the prey will only be able to move a certain distance and the shadower will only be able to see the prey from a certain distance. The premise of this work is that if the shadower can capitalise on such information it will be able to estimate the distance to the prey sufficiently accurately to perform motion camouflage. Therefore at each time step the shadower is provided two sensory inputs:

- the direction of the prey;
- the image motion of the prey (i.e. the angle subtended at the shadower's retina by the previous and present position of the prey, relative to the shadower at each respective time instant).

To make prediction possible the shadower is also afforded a short memory of recent prey image motion and in order to interpret this a memory of its own recent movement. The shadower is expected to estimate the distance to the fixed point based on this memory of its movement. Note that this task is not beyond the wandering spider *Cupiennius* which is capable of dead reckoning based entirely upon signals from mechanical stress sensors located on its legs [4].

The controller outputs the direction to move and the rotation about which to turn. The shadower is trained to rotate so as always to view the prey frontally and keep the fixed point directly behind. Assuming accurate operation this means that the direction of the fixed point is always known, leaving only the distance to be estimated. This rotation was suggested and observed in the hoverflies investigated in [1]. The shadower step size is held constant and modelled implicitly by the system.

B. Control system architecture

The control system is formed from three back-propagation artificial neural networks [5] (**Fig. 4**). The first network is trained to output the distance to the fixed point. The second to output the direction in which to move and the third the rotation about which to turn.

Before making a movement, the controller is provided with the direction to and image motion of the prey. This information, combined with:

- a memory of the prey image motion over the previous two time steps;
- a memory of the previous three movements made by the shadower;
- and the previous estimate of fixed point distance

is given to the distance network. This network estimates the distance to the fixed point. The new distance estimate is passed to the direction and rotation networks (and fed back to form part of the next input to the distance network). Otherwise their input is identical to that of the distance network. The direction and rotation of movement are then decided and output. These are also stored in memory and subsequently fed back as input to all networks.

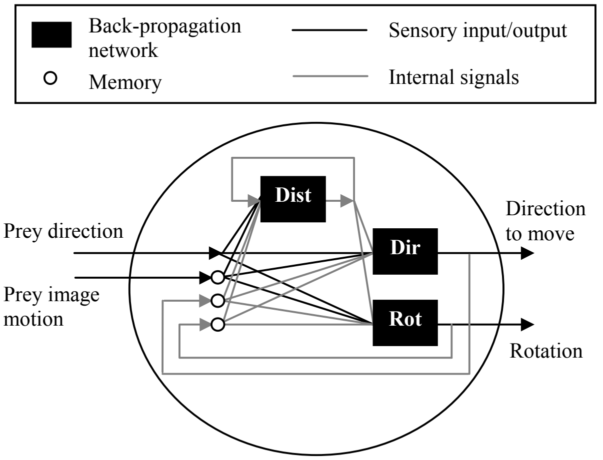

Fig. 4. Control system architecture. See text for explanation.

Each network is comprised of two hidden layers, the first consisting of 20 nodes, the second 10. The networks were trained with an adaption rate of 0.1 and a momentum term of 0.5. All control system parameters were selected on the basis of a series of preparatory experiments.

C. Training and testing procedure

Training and testing were undertaken by running sets of trials. In each trial the shadower would attempt to perform motion camouflage on a prey following a predetermined trajectory. Initial shadower positions were set at random bearings and random distances within a given range of the initial prey position. The fixed point was always taken to be the initial shadower position.

Target outputs in training were calculated *aposteriori* from exact distance information using trigonometry. The path of the shadower was determined by the controller output. **Fig. 5** shows the movements trained for a variety of shadower positions. If the future prey position was in between shadower and fixed point or within one step of the shadower, the shadower was trained to move straight towards this.

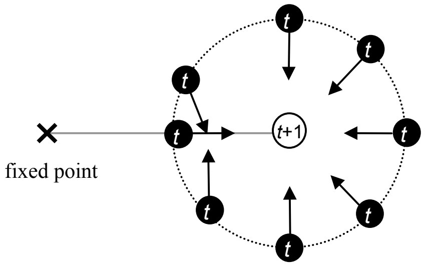

Fig. 5. Movement trained with Predictive training algorithm. The shadower is also trained to rotate to view the prey frontally.

```
if ( distance to constraint line <= step size )
        move to the intersection of step and constraint line nearest
        to the prey
else
        move directly towards the nearest point on the constraint line
```

In testing, for each shadower movement, recordings were made of the:

- *directional* and *rotational error*, the angle difference between both direction and rotation output from the targets;
- *visual error*, the angle difference of the shadower from the direction of the fixed point as seen from the prey (i.e. the angle subtended at the prey by the fixed point and the shadower).

These were used to quantify the quality of performance.

D. Prey trajectories

Two types of prey trajectory were used to train and test the networks. The first were automatically generated regular arcs. Here the prey moved a constant distance and changed direction by a constant angle within the range –0.04 to 0.04 radians at each time step. These trajectories were therefore totally predictable. In training the direction change of the prey was selected randomly at the start of each trial. In testing the direction changes were selected to span the range uniformly (i.e. if 5 trajectories were selected for testing the corresponding set of direction changes would be {-0.04, -0.02, 0, 0.02, 0.04}). To enable the shadower to always approach the prey, the distance moved by the prey at each time step was set to 9/10 that of the shadower.

The second set were the digitised flight paths of real hoverflies. Overall 91 sequences consisting of 40 frames were filmed. 75 of these sequences were randomly selected for training. The remaining 16 were used solely for testing. Again, to allow the shadower to always approach the prey, the hoverfly trajectories were pre-processed so that the mean step size in each trajectory was 9/10 that of the shadower.

The actual step sizes (relative to the distance separating shadower and prey at the start of the trial) were set so that there would be an obvious difference in visual error between implementations of the responsive and predictive algorithm. In explanation, in the simulation, shadower and prey move at the same time and the visual error is recorded after each movement. If the prey step size is very small there will be virtually no difference in visual error between the responsive algorithm and perfect camouflage and hence no advantage offered from prediction.

E. Experimental procedure

Beyond successfully training a controller, the following areas were identified as being of particular interest:

- Whether it is possible to improve motion camouflage through prediction.
- Whether the control systems are able to perform on the alternative trajectory type to those used in training.

Overall, 6 controllers were trained in the manner described in II.*C* to implement the predictive algorithm (parameters used in training and testing are displayed in **Table. 1**). Three were trained on the hoverfly training data and the remaining three on automatically generated regular trajectories.

Each control system was trained for a fixed number of iterations. During training, tests were performed and a record kept of the weight configuration of the best performing state of the controller. Specifically, trials were run from 10 randomly selected starting positions on ten different prey trajectories (in the training set). Accuracy was measured by the mean directional error. On completion of training, the best performing state of the controller found during training was selected as the final state (it had been found that the state of the control system after the last training iteration was not necessarily the best).

Following training, each control system was more extensively tested. 100 trials were run from randomly selected starting positions on each of 16 prey trajectories. For both trajectory types, the controllers with the least mean directional error were selected for further testing. These two also gave the least rotational and visual error means (respective to test trajectory type). These were then tested on the alternative trajectory type to those trained (i.e. the hoverfly trained controller was tested on regular trajectories and *vice versa*). Performance was compared to that obtained with the responsive algorithm (calculated using trigonometry from exact distances).

TABLE 1. Training and testing parameters

Number of steps moved per trial	39
Shadower step size (units)	5
Prey step size (units)	4.5
Initial distance range of shadower from prey	$200 \le d \le 400$
Initial bearing of shadower from prey (rad)	$0 \le b \le 2\pi$

III. RESULTS

Motion camouflage trajectories generated from the controllers are displayed in **Fig. 6**. Casual inspection suggests that the performance is very accurate, with the shadower able to successfully approach from different start positions. Camouflage is almost perfect on the regular trajectories, and although the shadower has more difficulty with the erratic movement of the hoverfly, it is able to keep to the constraint line remarkably well (especially considering that the hoverfly occasionally moves so fast that it is impossible for the shadower to actually reach the constraint line e.g. **Fig. 6b**).

Fig. 7 compares the **(a)** visual, **(b)** directional and **(c)** rotational errors of each controller on each test. As might be expected, for each test trajectory type the least errors were given by the controller trained on similar data (i.e. the controller trained on hoverfly trajectories performed better than anything else on the hoverfly test data). Notably, the errors were also less than those encountered from use of the responsive algorithm, indicating that prediction had improved camouflage in all respects. All differences mentioned in this section are supported statistically (**Table. 2**).

The other area of interest was how well the control systems performed on the alternative data type to that on which they had been trained. It was found that the hoverfly trained control system performed well on the regular test set. The performance was better in all respects than the responsive algorithm, indicating that the controller was capable of making accurate predictions. The opposite was not true. The control system trained on regular data performed poorly on the hoverfly trajectories and far less accurately than the responsive algorithm. The success of the hoverfly trained controller suggests that many of the predictable characteristics of the regular trajectories were also present in the hoverfly trajectories. Whether the poor performance of the regular control system on hoverfly data is indicative of a difference in dependency on inputs between controller types has yet to be ascertained. For instance, in the regular trajectories the prey never changed speed or abruptly changed direction. Subsequent camouflaged moves were therefore always very similar. It is possible that the shadower could have become too reliant on its previous moves to determine its next to reliably track the hoverfly. This awaits further investigation.

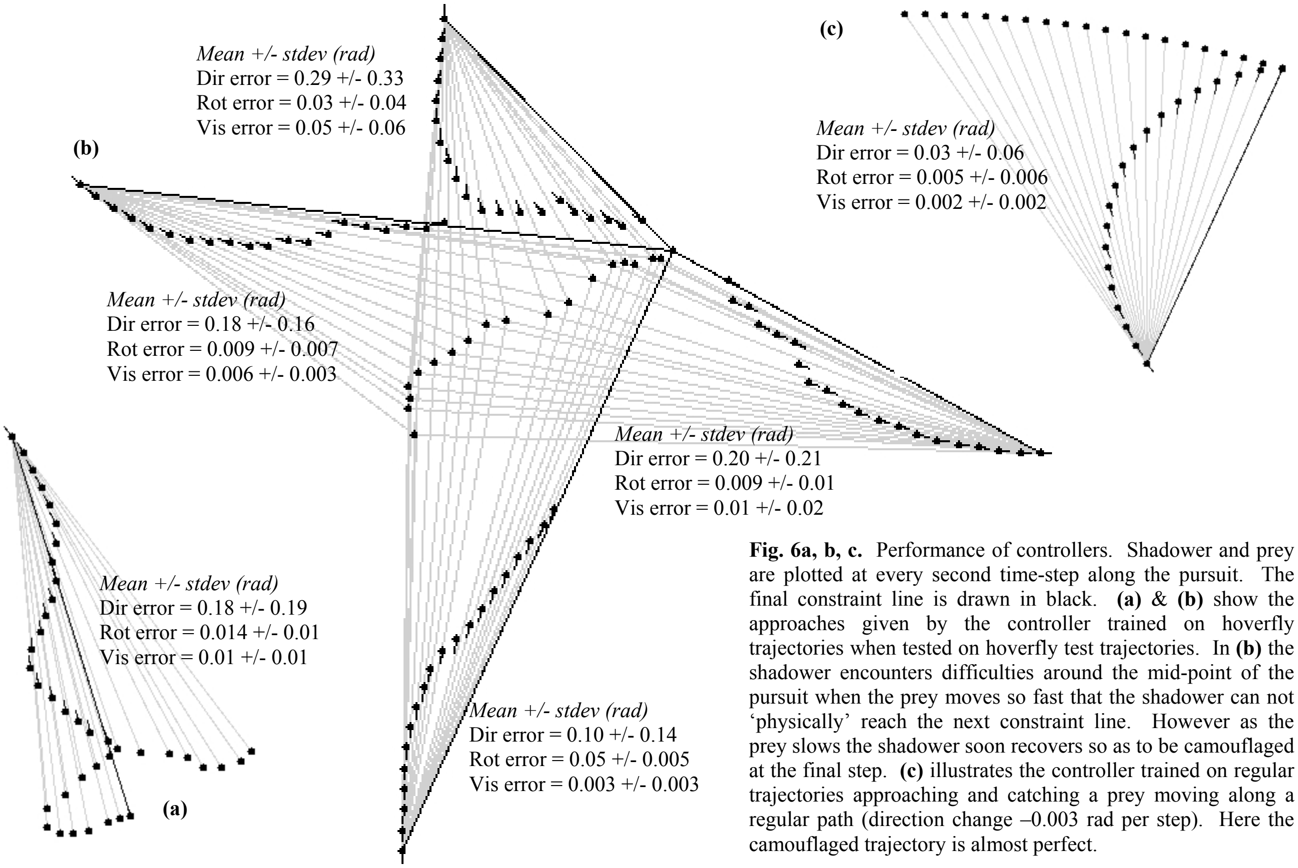

Fig. 6a, b, c. Performance of controllers. Shadower and prey are plotted at every second time-step along the pursuit. The final constraint line is drawn in black. **(a)** & **(b)** show the approaches given by the controller trained on hoverfly trajectories when tested on hoverfly test trajectories. In **(b)** the shadower encounters difficulties around the mid-point of the pursuit when the prey moves so fast that the shadower can not 'physically' reach the next constraint line. However as the prey slows the shadower soon recovers so as to be camouflaged at the final step. **(c)** illustrates the controller trained on regular trajectories approaching and catching a prey moving along a regular path (direction change –0.003 rad per step). Here the camouflaged trajectory is almost perfect.

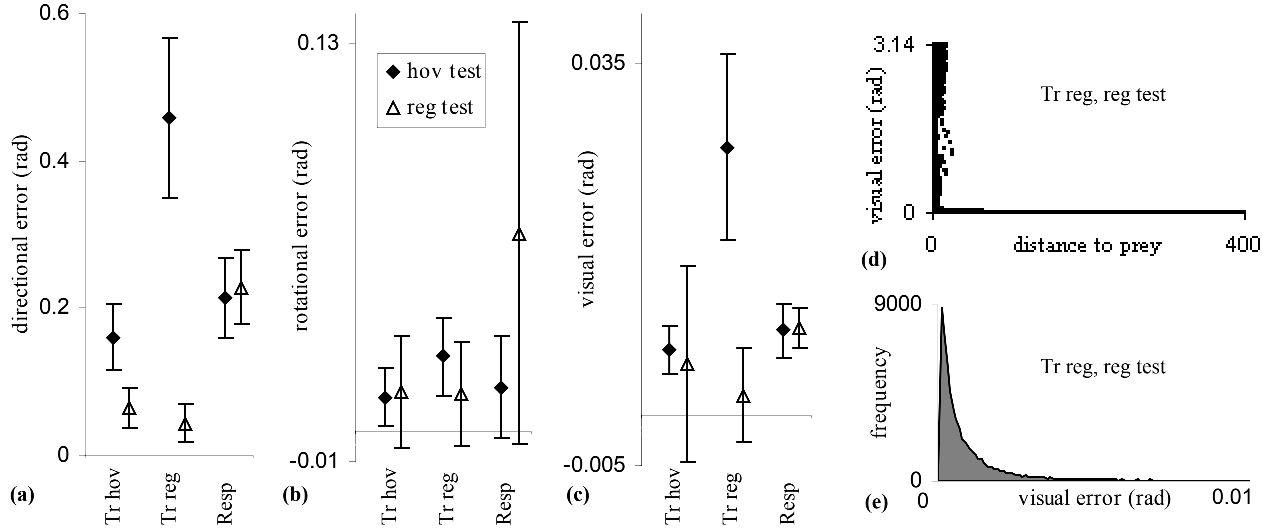

Fig. 7a-c. Comparison of the mean +/ stdev **(a)** directional, **(b)** rotational and **(c)** visual errors of the different controllers. Standard deviations are scaled by **(a)** 0.2, **(b)** 0.2, **(c)** 0.1 for the purposes of presentation. Each result is representative of 100 trials of 39 steps run against each of 16 prey trajectories. The high rotational errors of the responsive algorithm were incurred when the shadower was very close to the prey (and rotated to view the previous prey position). This happened more frequently with the regular trajectories when shadower and prey were more likely to move directly toward one another and meet. In **(c)** the visual error was not recorded after the shadower was ≤ 10 units distance from the prey. Beyond this distance, any slight directional error could incur a very large visual error, exerting a disproportionate influence on the mean and standard deviation. **(d)** shows the relationship between the distance separating shadower and prey and visual error for the controller trained on regular trajectories, tested on regular trajectories. As can be seen, there are only large errors when the shadower is close to the prey. **(e)** shows the frequency distribution of visual errors for **(d)**. The vast majority of errors are below 0.005 rad. 3427 (of a total of 62400) entries are above 0.01 rad and not shown on the plot.

TABLE 2 a, b, c. Scheirer-Ray-Hare (non-parametric equivalent of two-way ANOVA) tests investigating the difference in **(a)** directional, **(b)** rotational and **(c)** visual errors (recorded only when the shadower was at a distance ≥ 10 units from the prey. See **Fig 7** for explanation) between the different controllers (and responsive algorithm) and test trajectory types. Replicates were the mean error per trial (16 prey trajectories with 100 trials per prey trajectory). In each case highly significant differences were found between controllers and between test trajectory types. There was also a highly significant interaction, a consequence of the poor performance on hoverfly test data of the controller trained with regular trajectories, contrasting to the improved performance of the hoverfly trained controller on regular trajectories.

It was found through exhaustive use of Mann Whitney-tests (and z -tests in the cases of testing for a difference between each test trajectory type in directional and visual errors incurred by the responsive algorithm, where variances met test assumptions) that there were highly significant differences (all $P \leq 0.0001$) between each possible combination of main effects but one: There was no significant difference between test trajectory types in the visual error of the responsive algorithm ($P = 0.75$).

(a) Directional error

Source	*df*	*MS*	*H*	*P*
Contr	2	16.91	1134.42	< 0.0001
Test traj	1	64.78	2175.93	< 0.0001
Interaction	2	39.62	2657.98	< 0.0001
Residual	9594	0.011		
Total	9599	0.03		

(b) Rotational error

Source	*df*	*MS*	*H*	*P*
Contr	2	0.7	122.28	< 0.0001
Test traj	1	0.44	76.86	< 0.0001
Interaction	2	0.91	158.96	< 0.0001
Residual	9594	0.005		
Total	9599	0.006		

(c) Visual error

Source	*df*	*MS*	*H*	*P*
Contr	2	0.08	58.78	< 0.0001
Test traj	1	0.04	14.69	< 0.0001
Interaction	2	0.21	154.29	< 0.0001
Residual	9594	0.003		
Total	9599	0.003		

IV. DISCUSSION

Before training, the main uncertainty was whether it would be possible for the control systems to gain an adequate concept of the distance to the prey from their inputs to perform motion camouflage. Prey image motion (see II.*A* for definition) was expected to be of particular importance in this respect, certainly during the earlier stages of training. However as the shadower is trained to continually view the prey frontally, thus in effect attempting to minimise prey image motion, image motion will act more as a measure of error for an accurate shadower. If the shadower rotates correctly, the degree of rotation will provide similar information to that which would have been given by prey image motion should the shadower not have rotated. Regardless, without any other knowledge, prey image motion and past shadower rotation can only reveal whether the prey has moved to the left or right (assuming the change in shadower position is also known).

In practice, at large distances, simply knowing that the prey has moved to the left or right may be sufficient for camouflage. As the lateral component of the correct movement will always be relatively small, as long as the shadower moves in roughly the right direction, it is unlikely to be detected. As the shadower approaches, its movement must be increasingly accurate to maintain camouflage. Then, to be useful, prey image motion must be combined with additional information.

As the speed of the prey is constrained (unavoidably in the real world and as stated in the simulation), the shadower, in training, could gain an impression of the frequency distributions of prey image motions and of its own movements at different distances from the prey. Also it could learn something of the changes in these measurements from one step to the next. For example, geometry dictates that the range and rate of change of image motion and shadower rotation will in most cases be greater the closer the shadower is to the prey. It is thought that choice of movement must be based on this knowledge. For example, although ambiguous for absolute measurements, the magnitude of rotation and image motion could be used to determine the maximum possible distance to the prey. Then knowing the distance to the fixed point could prove valuable for more accurate estimation (e.g. if this distance is large, then the shadower is likely to be relatively close to the prey). An experimental analysis intended to identify the relative importance and relationships between inputs will form part of future work. This shall involve comparisons of the performance of control systems trained with missing inputs and of existing control systems with ablated inputs.

In summary, this paper has proposed, simulated and tested control systems for motion camouflage. It has been shown that a high level of accuracy can be achieved with simple inputs and that performance can be improved through prediction. This is encouraging for future work which shall concentrate on the development of a control system for an autonomous robot. A more difficult task as the system will then have to cope with the real world as seen through noisy sensors and the inevitable differences between intended and actual movements.

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] V. Bruce, P. R. Green and M. A Georgeson, Visual perception, physiology, psychology and ecology 3rd Ed, Psychology Press, 1996.

[3] K. Kral, "Side-to-side head movements to obtain motion depth cues: A short review of research on the praying mantis," *Behavioural Processes,* Vol. 43, pp. 71-77, 1998.

[4] E. A. Seyfarth, R. Hergenroder, H. Ebbes and F.G. Barth, "Idiothetic orientation of a wandering spider – compensation for detours and estimates of goal distance," *Behavioural Ecology and Sociobiology*, Vol. 11:2, pp. 139-148, 1982.

[5] D. E. Rumelhart, G. E. Hinton and R. J. Williams, "Learning internal representations by error propagation," *Parallel Distributed Processing*, Vol. 1, pp. 318-62, 1986.

This work was funded by an EPSRC grant.