THE ROYAL SOCIETY

Received 2 September 2002
Accepted 6 November 2002
Published online 16 January 2003

Model of a predatory stealth behaviour camouflaging motion

Andrew James Anderson* **and Peter William McOwan**

Department of Computer Science, Queen Mary, University of London, London E1 4NS, UK

A computational model of a stealth strategy inspired by the apparent mating tactics of male hoverflies is presented. The stealth strategy (motion camouflage) paradoxically allows a predator to approach a moving prey in such a way that it appears to be a stationary object. In the model, the predators are controlled by neural sensorimotor systems that base their decisions on realistic levels of input information. They are shown to be able to employ motion camouflage to approach prey that move along both real hoverfly flight paths and artificially generated flight paths. The camouflaged approaches made demonstrate that the control systems have an ability to predict future prey movements. This is illustrated using two- and three-dimensional simulations.

Keywords: motion camouflage; stealth; sensorimotor; neural network

1. INTRODUCTION

Motion camouflage is a stealth strategy that allows one moving body (a shadower) to camouflage its motion from another moving body (the prey). The technique was first suggested by Srinivasan & Davey (1995) who observed that, in mating, male hoverflies may move in a manner consistent with motion camouflage to track females.

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 (a fixed point). This is achieved by the shadower ensuring that it is always positioned directly in between this fixed point and the prey (figure 1). The line connecting fixed point and prey is referred to as the camouflage *constraint line*. If the shadower does this, the prey will perceive no lateral motion (retinal slip) as the shadower approaches, only the image of the shadower increasing in size. If the shadower is at sufficient distance from the prey the rate of change of image size will be negligible and the shadower's motion camouflaged (for example, Srinivasan & Davey (1995) assume the shadower to be at such a distance that it appears to the prey as a structureless dot).

As an example, 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 not be alerted to the approach of the shadower. Note that motion camouflage does not require the fixed point to be an existing landmark; it could equally well be the initial position of the shadower at the onset of the pursuit (in which case the prey would regard the shadower itself as a stationary landmark). Also, as discussed by Srinivasan & Davey (1995), that motion camouflage may be employed to retreat from the prey or to conceal movement to a particular location. This work concentrates on the shadower approaching the prey. It also assumes that the fixed point is located at the initial position of the shadower and does not require the fixed point to be visible (either because the fixed point is not an existing object or because the shadower cannot see directly behind itself).

Srinivasan & Davey (1995) showed that, in theory, if the shadower can accurately estimate its current distance from the fixed point ρ, it could camouflage its motion by: (i) always viewing the prey frontally; (ii) always pointing radially away from the fixed point (or estimated position of the fixed point as would be the case in this paper); and (iii) making corrective yaws of angle $\Delta\theta$ and the lateral component of its movement $\Delta\lambda$ such that the magnitudes are related by $\Delta\lambda/\Delta\theta$ (figure 2*a*).

However, this algorithm makes the assumption that the shadower already knows the appropriate yaw or lateral component of movement to make. Unless the shadower can be certain of the prey's pattern of movement in advance, neither $\Delta\theta$ or $\Delta\lambda$ can be known instantaneously and both must be estimated. In order to estimate either, the shadower must have a concept of its distance from the prey (figure 2*b*).

Given that the fixed point is directly behind the shadower, is at a known distance ρ and the current angle subtended at the fixed point by shadower and prey is $\Delta\theta$ ($\Delta\theta = 0$ if the shadower is correctly camouflaged) the shadower has the opportunity to calculate the appropriate move to camouflage itself perfectly for the time instant at which it commenced its move. But, on completion of its move, the prey is liable to have moved elsewhere. If the shadower does not consistently react sufficiently quickly to the prey's movement, the prey is likely to notice the lag in the shadower's response and consequently camouflage will be lost. To perfectly camouflage itself, the shadower must be able to predict future camouflaged positions, in doing so implicitly predicting future prey motion.

This paper addresses these issues and presents the first computational implementation of motion camouflage known to the authors. Autonomous shadowers are simulated employing predictive motion camouflage to approach prey moving along both the 2D digitized flight paths of real hoverflies and 2D and 3D artificial trajector-

*Author for correspondence (aja@dcs.qmul.ac.uk).

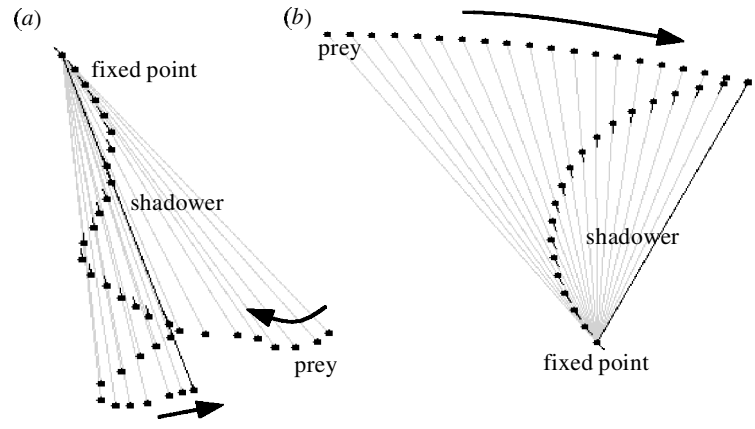

Figure 1. Camouflage trajectories showing the shadower approaching a prey moving along (*a*) the flight path of a real hoverfly (*b*) an arc. The shadower is depicted by the dots with tails, the prey, the dots without. At each instant the shadower is expected to lie on the camouflage constraint line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

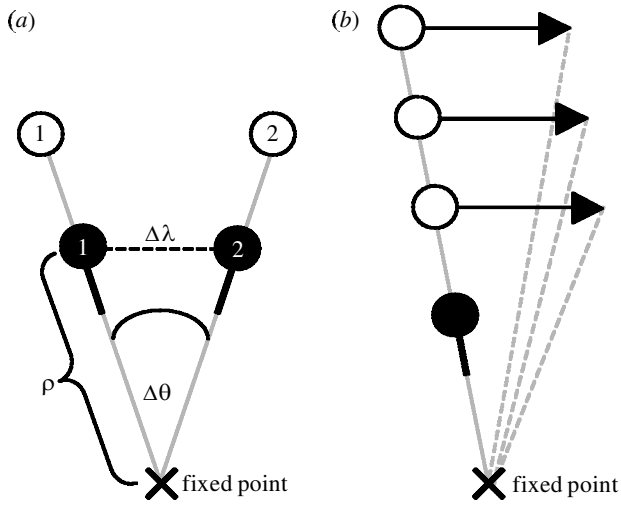

Figure 2. (*a*) Illustration of the components of the motion camouflage algorithm of Srinivasan & Davey (1995), $\Delta\lambda$ represents the lateral component of shadower movement. $\Delta\theta$ the yaw necessary for the shadower to face away from the fixed point, and ρ the distance from shadower to the fixed point. To be able to calculate any one of these terms exactly the shadower must know the other two. As illustrated in (*b*), to be able to estimate $\Delta\theta$ the shadower must have a concept of the distance to the prey. The same prey movement is shown with the prey (open circles) at different distances from the shadower (filled circles); each would require the shadower to move differently to maintain camouflage.

ies. The approaches demonstrate the ability of the shadowers' control system to estimate the position of the fixed point using dead reckoning. The shadowers are controlled by artificial neural networks operating with realistic levels of input information.

2. MATERIAL AND METHODS

(a) *Shadower control system design*

To be able to calculate a camouflaged approach, the shadower must estimate the distance and direction of the prey and fixed point. In designing the simulation, emphasis was placed on supplying the control systems with information that it is reasonable to assume could be retrieved by a real world shadower, be it biological or artificial.

The shadower is assumed to have rudimentary vision (see below for specifics). Spatial information could be acquired from other senses such as hearing (e.g. owls and, with echolocation, bats) and electromagnetic senses (e.g. dogfish); however, 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). As motion camouflage will be effective at distances where the prey cannot accurately judge the shadower's proximity, the amount of depth information that the shadower can extract from its own vision is likely to be limited. For example, binocular depth cues such as binocular stereopsis, convergence and accommodation (see Bruce *et al.* 1996) are likely to be useless. Therefore, to maximize the simplicity of the model, the only information that the shadower has provided by its vision is the position of the prey in its visual field. The remaining input is feedback of previous motor outputs. Specifically, the outputs are the direction in which to move that, it is estimated, will translate the shadower to a camouflaged position as close to the prey as possible. Second, the rotation to face away from the estimated position of the fixed point (i.e. $\Delta\theta$ in the 2D example in figure 2). The speed of the shadower is held constant. From the above information the control system is expected to estimate the current distance of the fixed point (the direction of the fixed point having already been estimated implicitly by the rotation output). As such, this is a case of *idiothetic* path integration (as opposed to *allothetic*, where the shadower would make use of external reference points; Papi 1992). In general, homing animals tend to combine allothetic and idiothetic path integration, the predominant method of navigation in bees and ants being allothetic based on a sky light compass (Wehner *et al.* 1996). Nevertheless the wandering spider, *Cupiennius*, is capable of dead reckoning based entirely upon mechanical stress sensors located on its legs (Seyfarth *et al.* 1982).

Notably, in biological neural systems there tend to be at least as many reciprocal connections as feed-forward. Even open-loop systems such as that controlling the vestibulo-ocular reflex (VOR) (responsible for instigating compensatory eye movements to stabilize the image on the retina with head movement) use parametric feedback for calibration (Churchland & Sejnowski 1992). However, feedback in the motion camouflage system presented bears more resemblance to the feedback loops suggested as a possible method of velocity storage in the VOR (Anastasio 1991).

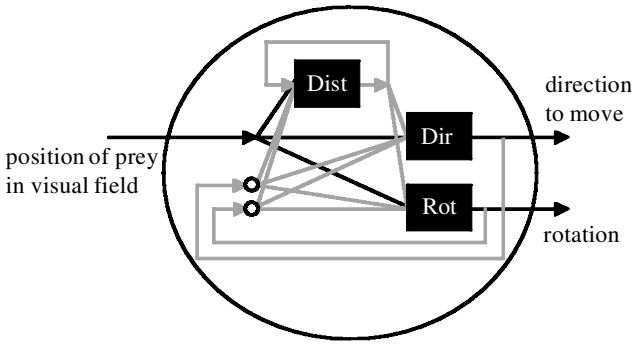

Figure 3. Control system architecture. 'Dist' indicates the network estimating fixed-point distance, 'Dir' indicates the network estimating movement direction and 'Rot' indicates the network estimating the rotation. See text for explanation. Key: black rectangles, multiplayer perceptron; open circles, memory; solid lines, sensory input/output; broken grey lines, internal signals.

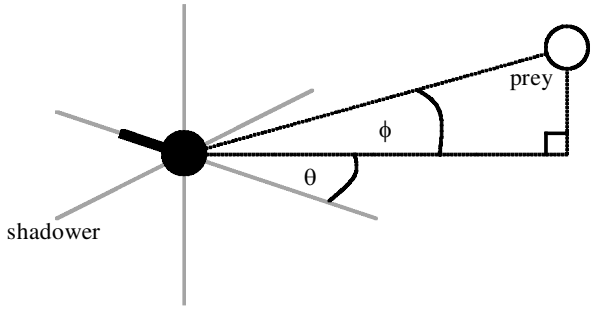

Figure 4. Representation of the position of the prey in the shadower's visual field, using two angles, the azimuth θ and elevation ϕ. Movement direction and rotation (both not shown) are similarly represented with two angles.

The architecture of the control system is shown in figure 3. The system is formed from three multilayer perceptron neural networks (for further information on all aspects of neural networks discussed here see Haykin 1999). These networks loosely resemble the structure of biological neural systems, in that they are formed from layers of connected artificial neurons, where the output of each neuron is a nonlinear function of its inputs that have been modified by synaptic weightings. The first network is trained to estimate the current distance from the fixed point; the second, the direction to move; and the third, the appropriate rotation. The position of the prey in the shadower's visual field, shadower movement direction and rotation are each represented by two angles in a spherical coordinate system as illustrated for the position of the prey in figure 4.

At each time-step, the system receives new sensory input (namely the two angles identifying the current position of the prey in the visual field) and the previous motor outputs are fed back into a memory that records the past three movement directions and rotations (for the purposes of prediction). Based on this information and its previous estimate, the distance network estimates the current distance to the fixed point. The new distance estimate is passed to the movement direction and rotation networks (and fed back to form part of the next input to the distance network). The movement direction and rotation networks make their respective estimations based on the new estimate of fixed-point distance, the current sensory input and the memory. These estimates are output from the control system and fed back to the memory.

Each of the networks is formed from two hidden layers, the first consisting of 40 artificial neurons, the second 25 neurons. Each artificial neuron uses the logistic activation function. Training was undertaken online using the standard back-propagation learning rule (Rumelhart *et al.* 1986), with a momentum term of 0.1 and a learning rate of 0.05 (all parameters were selected empirically). This learning rule brings the synaptic weights of the network from random starting values to values at which the overall network error is at a minimum. This is achieved through presentation of multiple example input and target output pairs to the network. With each presentation the synaptic weights are changed in proportion to the current error of the network (and learning rate and momentum term) to bring the overall output closer to the target (i.e. supervised learning).

(b) *Training and testing the control systems*

Training and testing were accomplished by running trials in which the shadower would attempt a motion-camouflaged approach on a prey moving along a predetermined trajectory. Initial shadower positions were selected randomly within a given distance range from the prey (between 200 and 400 distance units from the prey, with the shadower moving 5 units per step). At every step of the trial each network in the control system was trained to move to the closest camouflaged position to the prey. Target outputs were calculated a posteriori with knowledge of the exact locations of the shadower, prey and fixed point. In the instance that the prey was between the shadower and the fixed point, the shadower was trained to move directly towards the prey. If the constraint line was further than 5 units from the shadower, the shadower was trained to move as close to the constraint line as possible.

Two different measures of error were used to assess the performance of the control systems, these were recorded at every step along the pursuit: (i) visual error: the angle subtended at the prey by the centre of the shadower and the fixed point; (ii) direction error: the difference(s) between the control system's movement direction output and the target output. In the 3D simulations this was represented using two angles, the azimuth error and the elevation error (i.e. the angle difference between the actual azimuth output and the target azimuth output and the actual elevation output and the target). In 2D only the azimuth error was calculated.

(c) *Prey trajectories*

Control systems were trained and tested on prey moving along either real hoverfly flight paths (of species *Episyrphus balteatus*) or artificially generated trajectories.

Overall, 91 hoverfly flights were filmed, each consisting of 40 frames (24 frames s^{-1}). Seventy-five of these sequences were randomly selected for use in training the systems. The remaining 16 were used solely for testing. The sequences were filmed from plan view and consequently were represented in the model as points in a 2D plane. The digitized hoverfly trajectories were normalized for the simulation so that the mean distance covered per step by the hoverfly prey was 4.5 units (i.e. the average speed of the prey was nine-tenths that of the shadower, thus guaranteeing the shadower the chance to approach the prey).

Artificially generated trajectories were of two types: (i) regular; and (ii) stochastic. Both were restricted to 40 frames to match the length of the hoverfly sequences. During the regular trajectories, the prey was made to move a constant distance (4.5 units) and change direction by a constant angle at each step. As such, forthcoming movements were totally predictable. Regular

trajectories were represented in both two and three dimensions. In two dimensions the trajectories resembled regular arcs; in three dimensions the azimuth and elevation direction changes were independent (but constant throughout each trial). Direction changes were selected randomly from the range -0.04 to 0.04 rad before each trial.

The stochastic trajectories were generated as a 3D substitute for hoverfly trajectories (which could only be filmed in 2D). The change in movement direction at each step was controlled by a fraction of the previous direction change plus a random component. Step size was dependent upon the current direction change. Equations (2.1) and (2.2) formally define the direction change and step, where $\Delta\theta_t$ and $\Delta\phi_t$ are the change in azimuth and elevation, respectively (spherical coordinate system with prey as origin) at time t and r_t is the step size.

$$\Delta\theta_t = \frac{\Delta\theta_{t-1}}{p} + m \cdot z_t, \; \Delta\phi_t = \frac{\Delta\phi_{t-1}}{p} + m \cdot z_t, \tag{2.1}$$

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

$$r_t = 2.5 + 5 \times \left(\frac{m}{m + |\Delta\phi_t| + |\Delta\theta_t|} \right). \tag{2.2}$$

Thus r_t is limited to lie within the range 2.5–7.5 units, with a mean of *ca*. 4.5 units (in keeping with the other prey trajectories).

(d) *Experimental procedure*

Beyond the successful training of the control systems, the experimental procedure was designed to investigate the following two areas: (i) whether the control systems were able to make successful predictions of prey motion; (ii) whether the control systems would be able to approach prey moving in a different manner from that on which they had been trained (e.g. can a shadower that is used to tracking hoverflies successfully approach prey moving along a regular trajectory?).

In total, four control systems were trained, one for each prey trajectory type (2D-hoverfly, 2D-regular, 3D-regular, 3D-stochastic). Training followed the procedure described in § 2c and lasted for a fixed number of time-steps (10^8). During training, the accuracy of the systems was periodically tested and a record maintained of the best performing configuration of the control system seen so far. On completion of training this record was selected as the final state of the controller (the state of the controller at the final training step was not necessarily the best). The tests undertaken during training consisted of 10 trials run from each of 10 randomly selected starting positions on prey trajectories in the training set. Accuracy was measured by the mean azimuth movement direction error (see § 2b) over all 100 trials.

Following training, the control systems were tested more rigorously (as described in § 2c). The 2D control systems were tested approaching both hoverfly and 2D regular prey trajectories and the 3D control systems were tested on 3D regular and 3D stochastic prey trajectories (eight tests). The 2D tests consisted of 100 trials (run from randomly selected starting positions) on each of 16 prey trajectories. For the hoverfly prey, these trajectories were the previously unseen test set. For the regular prey, these were 16 trajectories uniformly spanning the direction change range (e.g. if there had been five trajectories the test set of direction changes would have been -0.04, -0.02, 0, 0.02 and 0.04 rad). The 3D control systems were tested on 100 trials on each of 100 prey trajectories. Here, all testing parameters were selected randomly.

In order to investigate for evidence of prediction, results were generated for each test corresponding to how a non-predictive shadower would have been expected to move (see § 1). Specifically, the non-predictive shadower moved so as always to be camouflaged according to the position of the prey at the previous time-step (i.e. at time t, the shadower would move to a position on the constraint line correct at $t - 1$ as close to the prey's position at $t - 1$ as possible). The movement of the non-predictive shadower was calculated using trigonometry from the exact coordinates of shadower, prey and fixed point.

3. RESULTS

Camouflaged trajectories generated by the different shadower control systems are shown in figure 1 (2D) and figure 5 (3D). In figure 1*a* the shadower is tracking a hoverfly trajectory, in figures 1*b* and 5*d* the prey trajectories are regular and in figure 5*a–c* they are stochastic. In all cases the shadowers can be seen to adopt an accurate camouflaged approach.

The results of the tests described in § 2d are displayed in figure 6. Two sets of analyses were carried out on these results. The first set investigated for evidence that the controllers were able to predict the prey's motion. The set of mean errors per trial (each trial lasting 40 time-steps) of the control systems tested on the same prey trajectory type to that with which they had been trained were compared with the corresponding set of mean errors given by the non-predictive shadowers using Mann–Whitney tests. There were 10 000 values of each error type (i.e. visual error, azimuth error and for the 3D simulations elevation error; see § 2b) for each experimental group in the 3D tests, corresponding to the 10 000 test trials (see § 2d), and likewise 1600 for each group in the 2D tests. Every test indicated a highly significant difference ($p < 0.0001$). Given that the average error of the control systems (when tested on familiar prey movement) was, with one exception, less than that given without prediction (see figure 6), this is thought sufficient evidence to demonstrate that the control systems were, in general, capable of predicting future camouflaged positions. The exception is the mean visual error of the 3D control system shown in figure 6*a*. Although the median visual error mean is less than that of the non-predictive trajectories, the mean is slightly greater. The differences in error between the control systems' approaches and non-predictive shadower's approaches were greater in the 2D simulations, especially for the predictable regular trajectories. The greater the s.e. of the 3D control systems' error in comparison with the non-predictive shadower, suggests that when the 3D control systems made errors, they were more serious than the non-predictive shadower. This is thought to be a consequence of occasional inaccurate predictions, and also indicative of the greater difficulty encountered tracking a prey moving in three dimensions.

The second set of analyses tested for differences in the control systems' accuracy when approaching different prey

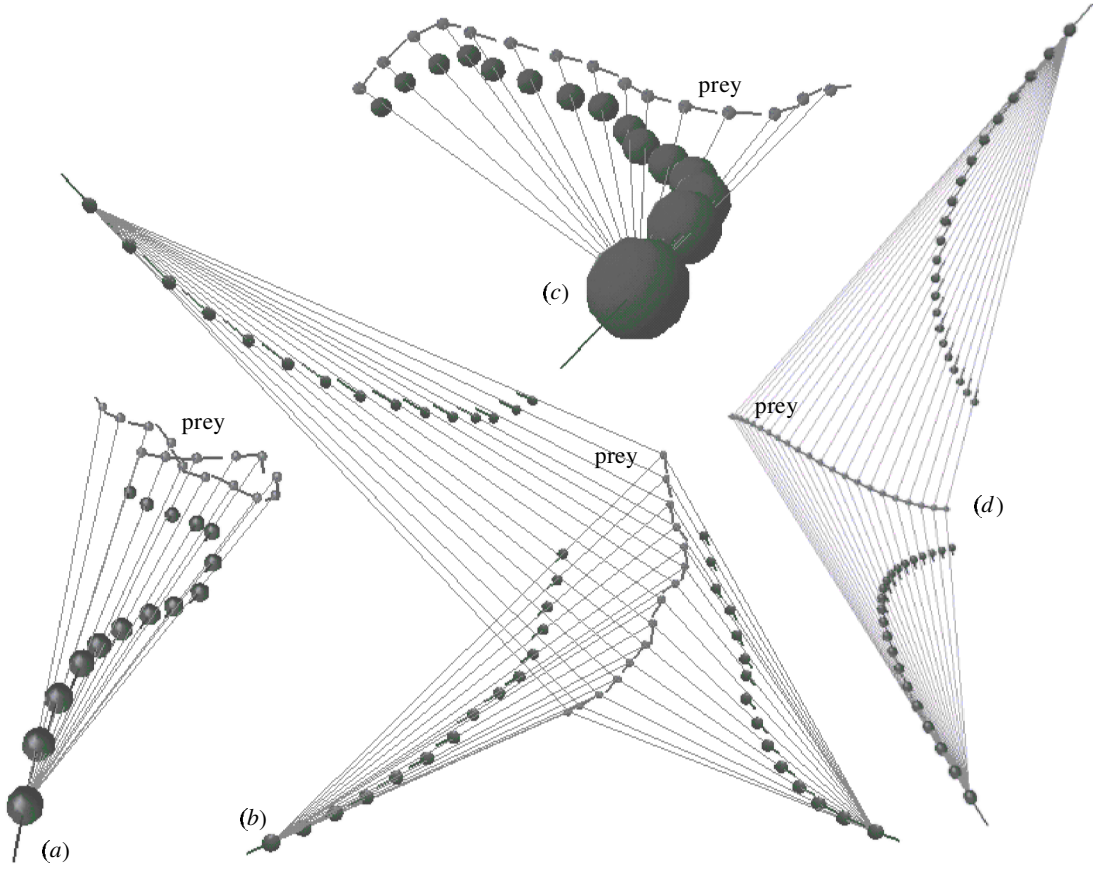

Figure 5. Example of camouflaged trajectories generated by the 3D shadower control systems shown from different viewpoints and rendered in 3D perspective: (*a–c*) show shadowers tracking prey moving stochastically (see § 2c); (*d*) shows two shadowers simultaneously approaching a prey moving in a regular pattern.

trajectory types. Scheirer–Ray–Hare tests (non-parametric equivalent of two-way ANOVAs based on ranked deviates) were used to compare the errors of the neural control systems when tested on the prey trajectory type with which they had been trained and the alternative. The main effects in the 3D tests were therefore 'stochastic versus regular test prey trajectories' (1 d.f.) cross classified against 'stochastic versus regular prey training trajectories' (1 d.f.). For the 2D tests, 'hoverfly versus regular test prey trajectories' (1 d.f.) was cross classified against 'hoverfly versus regular training prey trajectories' (1 d.f.). Replicates were the rank corresponding to the mean of each error type (see § 2b) per trial. This gave a total of 1600 replicates for each experimental group in the 2D tests (giving a total of 6399 d.f.) and 10 000 replicates for each group in the 3D tests (giving a total of 39 999 d.f.). For all error measures, all results (including interactions) were highly statistically significant ($p < 0.001$). In figure 6, it can be seen that each control system outperformed the others when tested on the prey trajectory type with which it had been trained. Also, camouflaged approaches were more accurate on the regular prey trajectories. The significant interactions were caused by the greater flexibility of the hoverfly and stochastically trained controllers which were able to track regular prey accurately, contrasting with the relative inability of the controllers trained on regular trajectories to track stochastic or hoverfly prey.

4. DISCUSSION

This paper has shown the simplicity of input information necessary to accomplish motion camouflage predictively, and also that the behaviour can be learnt by an artificial system that works in a manner reminiscent of a biological nervous system.

Regarding the task, what is especially interesting is that the control systems have been able to gain an adequate concept of prey depth from their inputs. There are a variety of clues that the control system has available for this. (i) It knows that it will be starting its approach from within a certain distance range of the prey (similarly it has been suggested that male hoverflies computing interception courses make the assumption that they will spot the female at a given distance; Collett & Land 1978). (ii) The further it estimates itself to be from the fixed point, the closer it is likely to be to the prey. (iii) Its previous rotation (combined with the feedback of its previous translation) may provide information in a manner similar to that in which peering enables animals such as praying mantids to estimate target distances. Praying mantids are able to accurately estimate the depth of stationary targets from the motion of the target's image encountered as they move their heads from side to side (Kral 1998). Ideally, as the shadower always faces its prey there will be no image motion. However, the corrective rotation necessary to

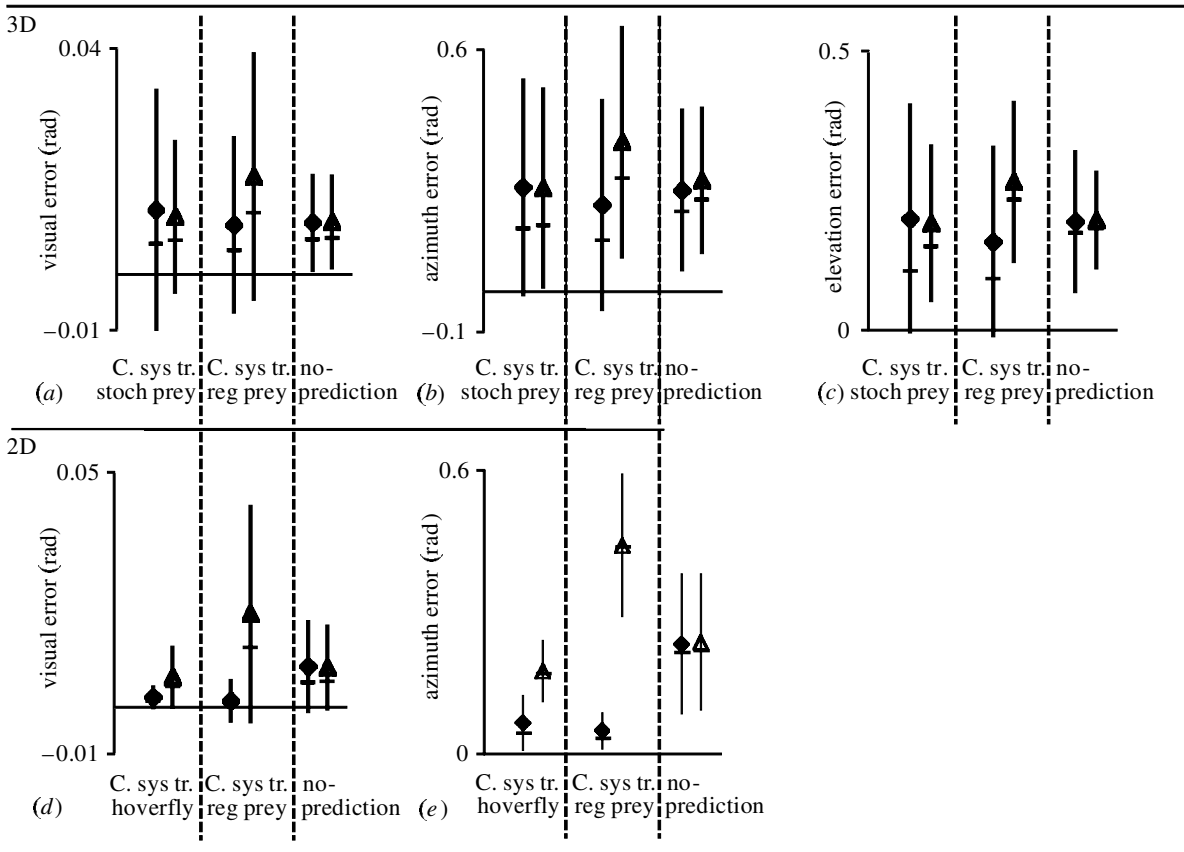

Figure 6. Comparison of the mean ± s.e. (and median) of the mean errors per trial resulting from the tests described in § 2d. The graphs in (*a–c*) display the results of the 3D tests, those in (*d*,*e*) the results of the 2D tests. Visual error is shown in (*a*) and (*d*). Visual error was recorded when the shadower was 15 distance units or more from the prey. Within this distance any slight directional error could incur a very large visual error, exerting a disproportionate influence on mean values. Movement direction error(s) (see § 2b) are shown in (*b*), (*c*) and (*e*). The type of prey trajectories (e.g. stochastic) used to train the control systems are indicated by the categories on the *x*-axis. Open or closed data points are used to indicate the prey trajectory type used for testing. Filled diamonds represent mean error when tested on stochastic/hoverfly prey; and horizontal lines represent the corresponding median error. Abbreviations: C. sys: control system; Tr: trained; Reg: regular prey movement; Stoch: stochastic prey movement.

keep the shadower facing the prey (as the shadower itself moves) will be identical to the image motion should the shadower not have rotated. In contrast to the praying mantis, whose target is normally stationary, both shadower and prey are moving, consequently making depth estimation more difficult. (iv) Finally, the shadower's errors may even be helpful in allowing it to infer its distance from the prey. The control system's only self-error measure is the extent to which the prey is non-central in its viewpoint (which is useful if the shadower is indeed facing away from the fixed point). Geometry dictates that the magnitude of this error measure will, on average, be greater for a similar movement error the closer the shadower is to the prey. Therefore, as a general rule, the greater the error, the closer the shadower is likely to be to the prey.

The fact that the control systems tended to specialize in tracking similar prey movement patterns to those on which they had been trained was not especially surprising. It is noteworthy that the control systems trained on the hoverfly trajectories should be able to generalize so well to tracking the artificial trajectories. This suggests that some of the predictable characteristics of the artificial trajectories were present (on occasion) in the hoverfly flight paths. The implications of this specialization are: (i) in nature, if motion camouflage shadowers do exist, they are likely to be specialist in certain prey; and (ii) artificially, motion camouflage control systems are likely to have to be tailored to particular prey. Conversely, returning to the point raised by Srinivasan & Davey (1995) of what counter strategies a prey could use to combat the shadower's camouflage, the prey may also need to be specialist in the avoidance of its potential predators. In investigating such areas computational modelling will prove an extremely valuable tool.

The authors thank two anonymous referees for their comments on an earlier manuscript. This work was funded by an EPSRC grant.

REFERENCES

Anastasio, T. J. 1991 Neural network models of velocity storage in the vestibulo-ocular reflex. *Biol. Cybern.* **64**, 187–196.

Bruce, V., Green, P. R. & Georgeson, M. A. 1996 *Visual perception, physiology, psychology and ecology*, 3rd edn. Hove, UK: Psychology Press.

Churchland, P. S. & Sejnowski, T. J. 1992 *The computational brain*. Cambridge, MA: MIT Press.

Collett, T. S. & Land, M. F. 1978 How hoverflies compute interception courses. *J. Comp. Physiol.* **125**, 191–204.

Haykin, S. 1999 *Neural networks a comprehensive foundation*, 2nd edn. Englewood Cliffs, NJ: Prentice-Hall.

Kral, K. 1998 Side-to-side head movements to obtain motion depth cues: a short review of research on the praying mantis. *Behav. Process.* **43**, 71–77.

Papi, F. (ed.) 1992 *Animal homing*. New York: Chapman & Hall.

Rumelhart, D. E., Hinton, G. E. & Williams, R. J. 1986 Learning internal representations by error propagation. In *Parallel distributed processing: explorations in the microstructure of cognition*, vol. 1 (ed. D. E. Rumelhart & J. L. McLelland), pp. 318–362. Cambridge, MA: MIT Press.

Seyfarth, E. A., Hergenroder, R., Ebbes, H. & Barth, F. G. 1982 Idiothetic orientation of a wandering spider—compensation for detours and estimates of goal distance. *Behav. Ecol. Sociobiol.* **11**, 139–148.

Srinivasan, M. V. & Davey, M. 1995 Strategies for active camouflage of motion. *Proc. R. Soc. Lond.* B **259**, 19–25.

Wehner, R., Michel, B. & Antonsen, P. 1996 Visual navigation in insects: coupling of egocentric and geocentric information. *J. Exp. Biol.* **199**, 129–140.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.