Experiments in automatic flock control

Richard Vaughan[⋆1,2], **Neil Sumpter**[1], **Andy Frost**[1] & **Stephen Cameron**[2]

[1]BioEngineering Division, Silsoe Research Institute, Silsoe, Bedford MK45 4HS, UK.
[2]Oxford University Computing Laboratory, Parks Road, Oxford OX1 3QD, UK.

Abstract. The Robot Sheepdog Project has developed a mobile robot that gathers a flock of ducks and manoeuvres them safely to a specified goal position. This is the first example of a robot system that exploits and controls an animal's behaviour to achieve a useful task. A potential-field model of flocking behaviour was constructed and used to invesigate methods for generalised flock control. One possible algorithm is described and demonstrated to work both in simulation and in the real world.

1 Introduction

The Robot Sheepdog Project has demonstrated a robot system that gathers a flock of ducks and manoeuvres them safely to a pre-determined goal position. No other robot system controls the behaviour of an animal and there existed no methodology for designing one. This work establishes such a methodology. The main research and methodological goal was to develop a machine that could usefully interact with an animal *without using the animal directly in the development process.*

The RSP is a collaborative, multidisciplinary project covering robot building, machine vision, behavioural modelling and ethological experiments. For an overview of the project as a whole, see [VHS97]. Work on machine vision has been reported as [SBT97] and [SB98].

The sheepdog's gather-and-fetch task was chosen because of its familiarity and the strong interaction between the dog, shepherd and flock animals. Using ducks instead of sheep allows us to experiment on a conveniently small scale, in a controlled indoor environment. Duck flocking behaviour is recognised by shepherds as similar to sheep; ducks are often used to train sheepdogs because of their relatively slow movement.

In order to identify the appropriate robot-animal interactions we built a minimal generalised model of the underlying flock behaviour. The hypothesis is that if the model accurately captures the basis of the behaviour, then a system which controls the model should control the real-world behaviour.

Models of flocking behaviour exist in the literature and are generally derived from Hamilton's observation that flocking may be produced by the mass action of individual animals, each seeking the proximity of its nearest neighbours [Ham71]. It was later suggested that this behaviour can be well modelled by an attractive 'force' acting between the animals, with the magnitude of the attraction varying with the inverse square of the animals' mutual distance [Par82] [WL91]. It is argued that this relationship represents a linear response to sensory information which itself varies with the inverse square of distance. Similar models have produced realistic computer animations of bird flocks [Rey87]. Flocks of mobile robots have also been demonstrated [Mat96].

These ideas are familiar in robotics, where such *potential field* techniques are used for navigation [CP94, Ch.10-11]. This class of algorithm uses the analogy of forces acting on particles, such that the robot will move as if it were a particle attracted or repelled from features in its environment. A robot is typically attracted to a goal position and repelled from obstacles.

The commonality of these animal and robot behaviour models forms the basis of an effective flock-gathering strategy, described below.

2 Rover the Robot Sheepdog

The experimental system comprises a robot vehicle, a workstation and a video camera (Figure 1, left). The vehicle was designed to work in a duck's environment: outdoors, on short grass, and in real time. Thus our

⋆ email: vaughan@comlab.ox.ac.uk

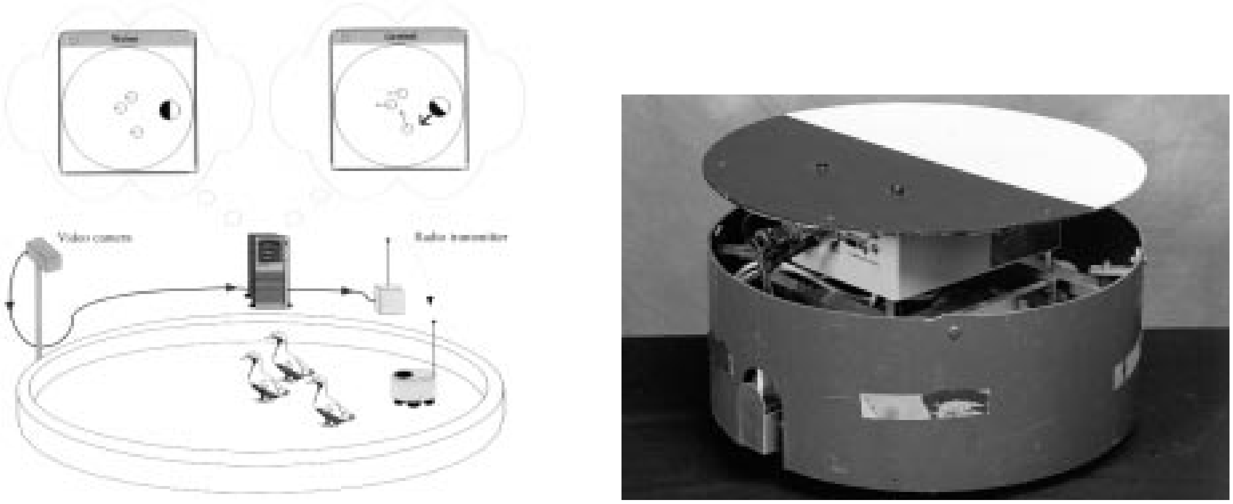

Fig. 1.: Robot Sheepdog system overview (left) and vehicle (right)

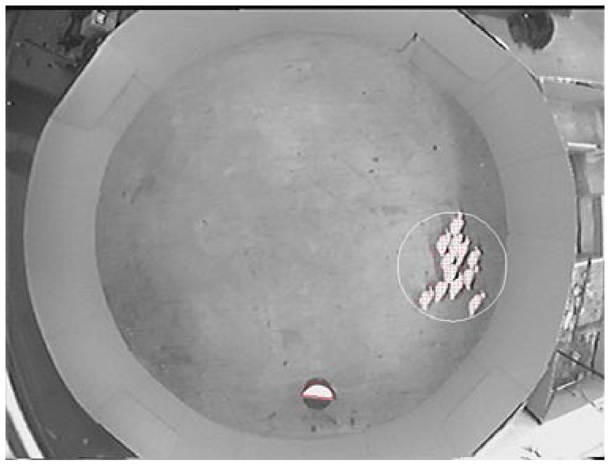

Fig. 2.: Example overhead camera image with the positions of the robot and flock overlaid by the tracker.

robot has an acceleration $\approx 1ms^{-2}$ and a top speed $\approx 4ms^{-1}$, which is about twice as fast as the ducks. It is covered in a soft plastic bumper mounted on rubber springs, ensuring duck safety. In the tradition of mobile robotics, we call it 'Rover' (Figure 1, right).

The vehicle and ducks are free to move in a visually uniform arena of 7m diameter, in view of the overhead camera. The positions of the robot and flock are determined by processing the video image stream. The robot's position and orientation are found by matching a template of its black and white cover to a region of the image. Tracking the ducks was a more unusual vision task. An ideal system would track the positions of individual ducks. It was concluded early on that there were no reliable, fast methods available to achieve this, certainly given our modest resources. However, it seemed likely that we could track the whole flock as one object, with some measurement of its size and shape. Such 'blob detectors' are common in machine vision, and are implemented via textbook techniques such as background subtraction and thresholding (see, for example, [BHS93]). Flock position would be defined as the as the center of area of the detected flock 'blob'.

This gave an interesting constraint to the rest of the system; it would have to work without knowing the positions of individual birds, but only with a center position, size and shape. In fact, the flock control algorithms that were devised do not require the shape information, so it was possible to abandon the shape-finding and produce a very fast tracker that finds just the center and radius of the flock. The final vision

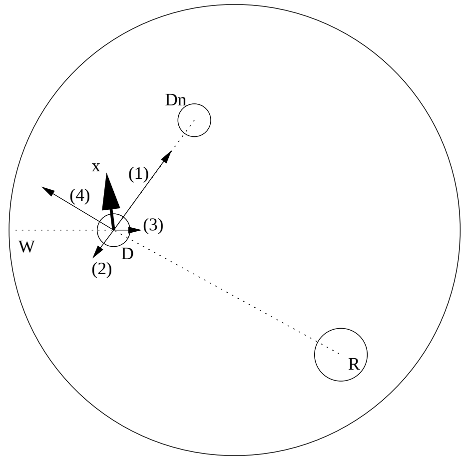

$$\mathbf{d} = \sum_{n=1}^{N} \left(\left(\frac{K_{f1}}{|\mathbf{DD_n}|^2} \right) \widehat{\mathbf{DD_n}} - L - \left(\frac{K_{f2}}{|\mathbf{DD_n}|^2} \right) \widehat{\mathbf{DD_n}} \right) - \left(\frac{K_{f3}}{|\mathbf{DW}|^2} \right) \widehat{\mathbf{DW}} - \left(\frac{K_{f4}}{|\mathbf{DR}|^2} \right) \widehat{\mathbf{DR}}$$

(1) (2) (3) (4)

Fig. 3.: Flock model (schematic not drawn to scale). Key: gain parameters $K_{1\to 4}$; repulsion bias parameter L (ensures repulsion > attraction at small distances, preventing collisions); ducklet position D, other ducklet D_n; Robot position R; Nearest point on wall W; algorithm terms ($1 \to 4$) and resultant $\mathbf{d}$ (where $\widehat{a}$ is the unit vector of $\mathbf{a}$).

system runs very quickly (update frequency > 25Hz), and has proved adequate for these experiments. Figure 2 shows a example image with the robot and flock correctly identified.

The robot's movement is guided by a flock-control algorithm running on the workstation. This algorithm takes the vision data (positions of the robot R, flock F and goal G) as input and returns a desired vehicle trajectory $(R, F, G) \to \mathbf{r}$. This is passed to the robot by radio modem, and a conventional high-speed proportional controller governs the robot's wheel speeds to approximate this path.

3 A model flock

A minimal simulation model of the duck-herding scenario was created, in which a flock of model ducks (ducklets) moves in a circular arena containing a model robot.

Given a ducklet's position D, the positions of the N other ducklets $D_{1\to N}$, the robot's position R and the nearest point on the wall W, the ducklet's movement vector $\mathbf{d}$ is determined by the function shown in Figure 3. The ducklets are (1) attracted to each other, aggregating the flock; (2) repelled from each other, preventing collisions and maintaining inter-ducklet spacing; (3) repelled from the arena wall, preventing collisions. A further term (4) which produces repulsion from the robot is proposed to model the aversive response of the ducklets to the robot. All these forces are scaled according to the inverse square of distance, and each ducklet moves according to the resultant of the forces acting upon it. The simulation produces a realistic-looking flock which can be manipulated by steering the model robot.

Note that the model describes a small subset of the ducks' behaviour. Of course, many other mechanisms generate the behaviour of real ducks, but our hypothesis is that this model captures enough of the real animals' behaviour to be a useful design tool. The model is a *generalised* description of flocking behaviour and as such could be applied to any flocking animal in two or three dimensions.

4 Flock control

Experiments with the simulator guided the development of two novel flock control algorithms which are closely related to the flock model described above. The most successful of these is presented here (see [VSFC98] for side-by-side comparison of both methods).

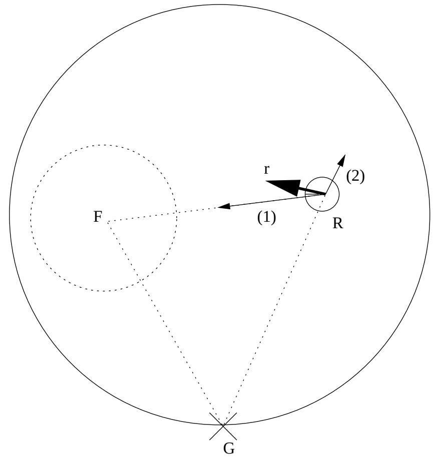

$$\mathbf{r} = (K_{r1}|\mathbf{GF}|)\widehat{\mathbf{RF}} - K_{r2}\widehat{\mathbf{RG}}$$
$$(1) \qquad\qquad (2)$$

Fig. 4.: Robot controller(schematic not drawn to scale). Key: gain parameters $K_{1,2}$; flock center F; Robot position R; Flock goal position G; algorithm terms ($1 \rightarrow 3$) and resultant $\mathbf{r}$ (where $\widehat{a}$ is the unit vector of $\mathbf{a}$)

The distance $|GF|$ in Figure 4 is the system variable we are trying to control, ie. reduce to zero. In a classical proportional controller a control output would be applied to correct this variable, with a magnitude proportional to the size of the error. If we include this term in the flock controller, we can design an analogous system whereby the repelling stimulus experienced by the ducks is proportional to their distance from the goal.

The robot's movement vector $\mathbf{r}$ is given by the function shown in Figure 4. The robot is (1) attracted to the flock with magnitude proportional to the distance from the flock to the goal; (2) repelled from the goal with constant magnitude.

4.1 Performance in simulation

The algorithm is first tested in simulation. A point on the arena boundary is chosen as the flock goal, 12 ducklets are placed randomly in the arena, and the robot positioned near the goal. The simulation starts and the positions of the robot and flock center are recorded for the next 3 minutes, as the robot attempts to manoeuvre the flock to the goal. This experiment was repeated nine times with the ducklets at different random start positions, and the robot at a slightly different position near the flock goal in each trial.

The results show that this controller performs the required task with some success. Figure 6 (A) shows a representative plot of the simulated robot and flock paths around the arena, while Figure 5 shows a series of screenshots from a similar trial. It can be seen that the flock is brought near the goal. The success plot Figure 6 (B) shows the distance of the flock to the goal over the length of the trial, plus the average distance over the entire trial. This is used as a measure of the the trial's success for comparision with other experiments. It can be seen that the flock-to-goal distance decreases rapidly then stabilises as the ducks settle near the goal. This trial scores an average flock-to-goal distance of 1.9m. The average score over all 9 trials was 1.8m.

4.2 Performance in real world

A similar experiment was then performed in the real world. A random point along the arena boundary is chosen as the flock goal. With the robot inactive and positioned near the goal, a flock of 12 ducks is introduced into the arena. After 3 minutes accomodation time, the robot is activated. The positions of the robot and flock center are recorded for the next 3 minutes, as the robot attempts to manoeuvre the flock to the goal. At the end of the trial, the robot is deactivated and the ducks move freely again for 2 minutes before being allowed out of the arena. This experiment was repeated three times with each of three flocks, with the robot at a slightly different position near the flock goal in each trial. Multiple flocks were used to

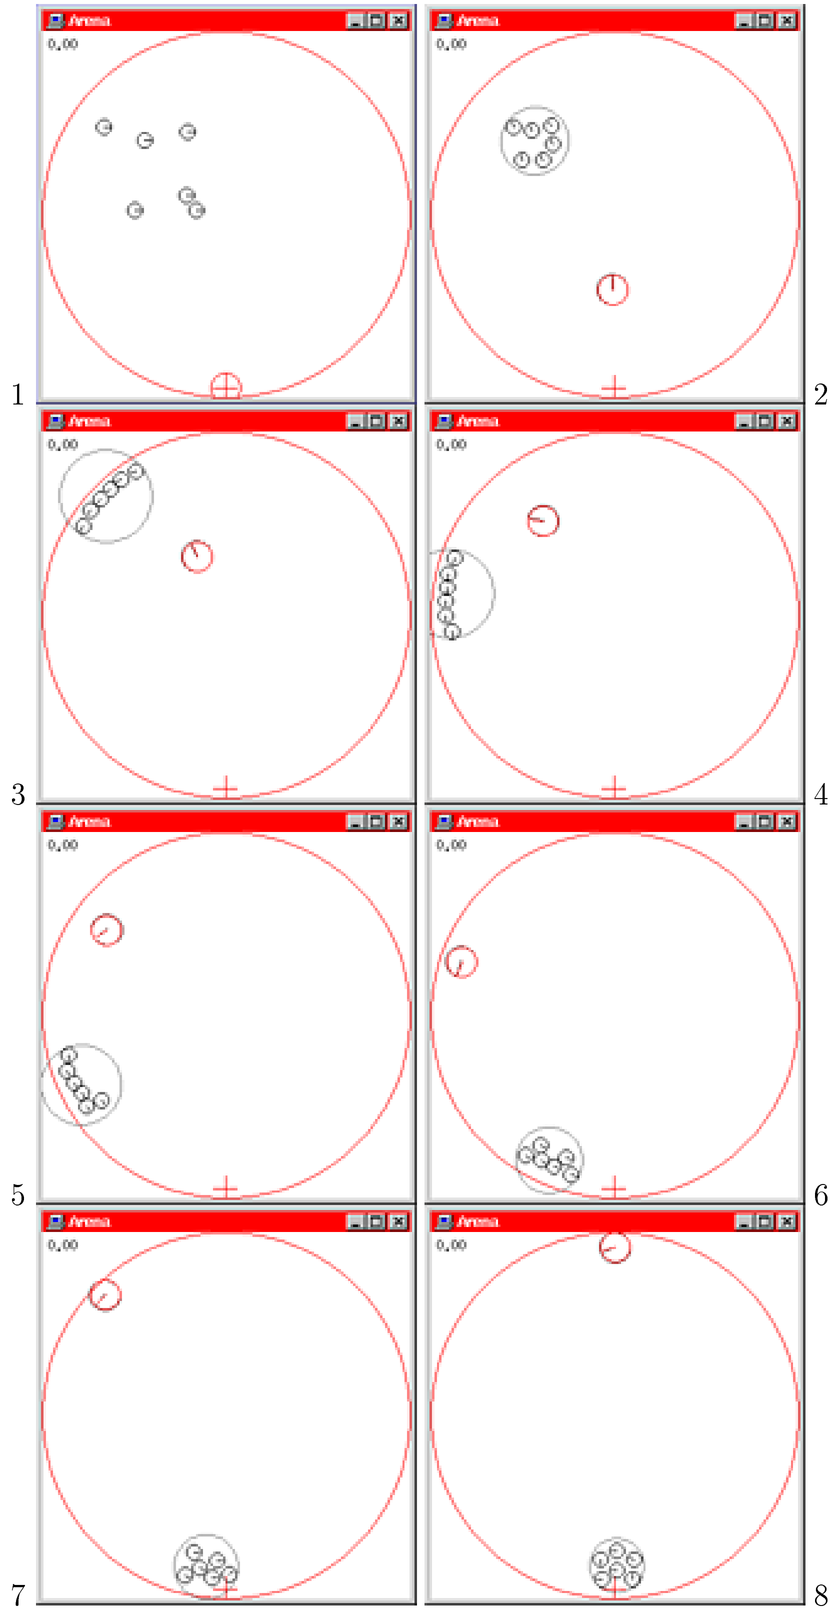

Key: ○ = robot, ○ = ducklets, + = flock goal

Fig. 5.: Sequence of images from the simulator during a trial, showing successful behaviour.

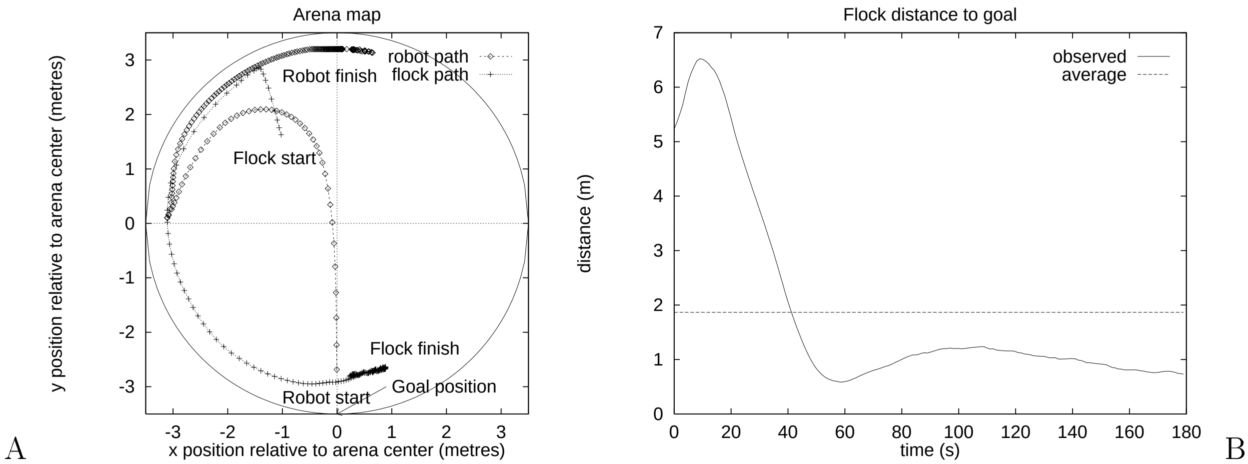

Fig. 6.: Simulation results.

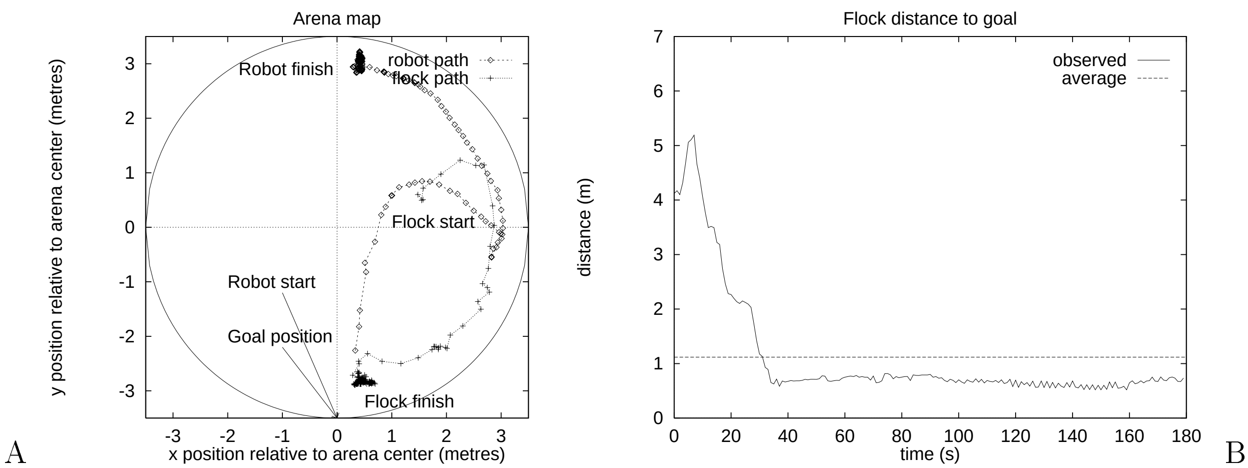

Fig. 7.: Real-world results.

increase the chance of variability in behaviour between trials. All the ducks were the same age and had been raised under similar conditions.

Figure 7 (A) shows a representative plot of the real robot and flock paths around the arena, while Figure 8 shoes a series of overhead camera images from the same trial. It can be seen that the robot approaches the flock, moving round behind them with respect to the goal. The flock moves away from the robot and towards the goal. As the flock approaches the goal, the robot is is less attracted to the flock and the goal repulsion becomes dominant. The robot retreats to the far side of the arena, applying minimum stimulus to the ducks. The ducks settle near the goal position. The success plot (Figure 7 (B)) clearly shows the initial fetching phase, followed by the stable, settled behaviour. This trial scores an average flock-to-goal distance of 1.1m. The average score over all 9 trials was 2.3m.

The larger average score was due to an overshoot effect, whereby the flock approached the goal but went past it and had to be fetched back by the robot. This effect is visible in the simulated trial in Figure 6 (the second, wider peak in plot B). The overshoot is caused by moving the ducks too quickly to the goal and not backing away quickly enough. Subsequent trials (simulated and real) have shown that the overshoot can be eliminated by tuning the gain parameter K_1 which controls the amount of attraction to the flock. The optimum setting of this parameter varies from flock to flock, and from day to day. As the success of this

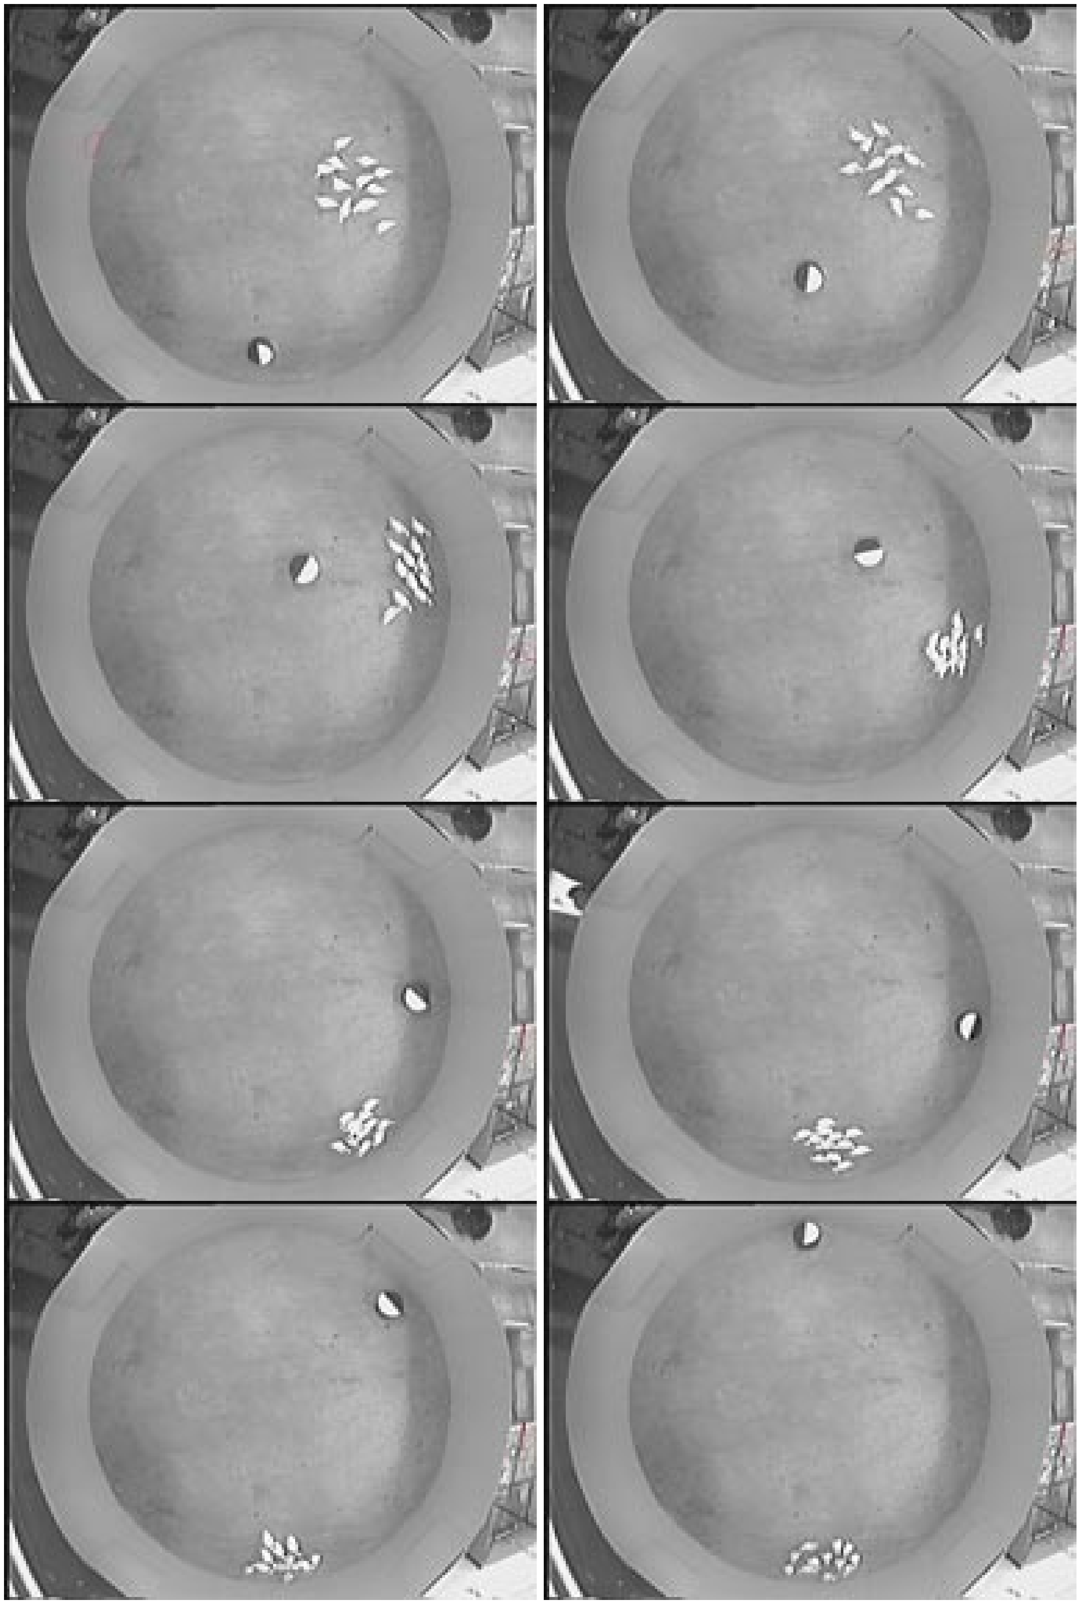

1 2 3 4 5 6 7 8

Fig. 8.: Sequence of images from the overhead camera during an experiment, showing successful behaviour. The goal position is at the bottom of the picture.

method varies (slightly) with this setting, work continues to find algorithms that may be more robust with respect to the inevitable variation between flocks.

5 Conclusion

We have demonstrated a robot system that achieves a sheepdog-like task, gathering and fetching live animals to a pre-defined goal position. We believe this is the first automatic system to exploit an animal's behaviour to achieve a useful task. A flock control method was designed and tested using a minimal simulation model of the ducks' flocking behaviour, and successfully transferred directly to the real world. We assert that the effectiveness of the simple method described is due to its close relationship to the mechanisms underlying flocking behaviour itself, and conclude (1) that behavioural simulations can be plausible engineering design tools, and (2) that such a methodology is appropriate for future animal-interactive robotics experiments.

6 Further work

This system uses a bird's-eye view of the arena, which would not be available to a real sheepdog or a robot with only on-board sensors. This arrangement was chosen at the start of the project for pragmatic reasons, to make it likely that we could produce a working system in the three-year life of the project. Sophisticated sensor engineering and processing was not a goal of the project: the focus has been on designing the appropriate interaction between the vehicle and the ducks.

However, the simulation model permits the investigation of deictic sensor modes without having to physically engineer them. We have devised alternative similar flock control algorithms which use only local sensing (ie. range data and/or vision) to achieve the same task. Results from this work will be presented elsewhere.

Acknowledgements

Thanks to Steve Crook, Paul Griffith, John Lowe and Robert Smith for their work on the robot, and to three anonymous reviewers for their helpful suggestions on this paper. The Robot Sheepdog Project is funded by the BBSRC and Silsoe Research Institute.

References

[BHS93] R. Boyle, V. Hlavac, and M. Sonka. *Image Processing, Analysis and Machine Vision.* Chapman and Hall, 1993.

[CP94] S. Cameron and P. Probert. *Advanced Guided Vehicles: Aspects of the Oxford AGV project.* World Scientific, 1994.

[Ham71] W.D. Hamilton. Geometry for the selfish herd. *Journal of Theoretical Biology*, 31:295–311, 1971.

[Mat96] M.J. Mataric. Designing and understanding adpative group behavior. *Adaptive Behavior*, 4(1):51–80, 1996.

[Par82] B. L. Partridge. The structure and function of fish schools. *Scientific American*, pages 114–123, June 1982.

[Rey87] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioural model. *Computer Graphics*, 21(4):25–34, jul 1987.

[SB98] N. Sumpter and A. Bulpitt. Learning spatio-temporal patterns for predicting object behaviour. Technical report, Univeristy of Leeds, School of Computer Studies, The University of Leeds, UK, 1998.

[SBT97] N. Sumpter, R.D. Boyle, and R.D. Tillett. Modelling collective animal behaviour using extended point distribution models. In *Proceedings of the British Machine Vision Conference*, 1997.

[VHS97] R. Vaughan, J. Henderson, and N. Sumpter. Introducing the robot sheepdog project. In *Proceedings of the International Workshop on Robotics and Automated Machinery for BioProductions*, 1997.

[VSFC98] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot control of animal flocks. In *Proceedings of the Joint Conference on the Science and Technology of Intelligent Systems (ISIC/CIRA/ISAS'98).* IEEE, in press, to appear 1998.

[WL91] K. Warburton and J. Lazarus. Tendency-distance models of social cohesion in animal groups. *Journal Of Theoretical Biology*, 150:473–488, 1991.

This article was processed using the TEX macro package with SIRS98 style