

Experiments in automatic flock control

Richard Vaughan^{a,b,1} Neil Sumpter^a Jane Henderson^a
Andy Frost^a Stephen Cameron^b

^a*BioEngineering Division, Silsoe Research Institute, Silsoe MK45 4HS, UK.*

^b*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 investigate methods for generalised flock control. One possible algorithm is described and demonstrated to work both in simulation and in the real world.

Key words: robot sheepdog; flock control; flock model; animal interaction

1 Introduction

The Robot Sheepdog Project (RSP) has demonstrated a robot system that gathers a flock of ducks in a circular arena 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 modeling and ethological experiments. An early overview of the project as a whole is given in [11] and results from our robot experiments have been presented at several conferences [12–14]. Our work on image processing and machine vision has been reported as [8–10].

¹ now at the Robotics Research Laboratory, University Of Southern California, Los Angeles, CA 90089-0781, USA. Email: vaughan@robotics.usc.edu




Fig. 1. A young sheepdog in training with a group of ducks. Photographed by the first author in Lancashire, November 1995.

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. For similar reasons ducks are often used to train sheepdogs, as seen in Figure 1.

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 [3]. 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 [6] [15]. 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 [7]. Flocks of mobile robots have also been demonstrated [5].

These ideas are familiar in robotics, where such *potential field* techniques are used for navigation [2, 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




Fig. 2. Robot Sheepdog system overview (left) and vehicle (right)



Fig. 3. The experimental arena (left) and view from overhead camera (right). Note the positions of the robot and flock overlaid by the tracker

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 2, left). The vehicle was designed to work in a duck's environment: outdoors, on short grass, and in real time. Thus our robot has an acceleration $\approx 1\text{ms}^{-2}$ and a top speed $\approx 4\text{ms}^{-1}$, 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 2, right).

The vehicle and ducks are free to move in a visually uniform arena of 7m diameter, in view of the overhead camera. The arena is shown in Figure 3 (left). 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 and although 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. 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 standard techniques such as background subtraction and thresholding (see, for example, [1]). Flock position was defined as the centre 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 centre 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 centre and radius of the flock. The final vision system runs very quickly (update frequency $> 25\text{Hz}$), and has proved adequate for these experiments. Figure 3 (right) 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 $(R, F, G) \rightarrow \vec{r}$. This is passed to the robot by radio modem, and a conventional high-frequency proportional controller governs the robot’s wheel speeds to closely approximate this vector.

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.

A potential field algorithm is used to generate movement for each ducklet. Given a ducklet’s position D , the positions of the N other ducklets $D_{1 \rightarrow N}$, the robot’s position R and the nearest point on the wall W , the ducklet’s movement vector \vec{d} is determined by the function shown in Figure 4. 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)



$$\vec{d} = \sum_{n=1}^N \left(\left(\frac{K_1}{(|\vec{D}\vec{D}_n| + L)^2} \right) \widehat{D}\vec{D}_n - \left(\frac{K_2}{|\vec{D}\vec{D}_n|^2} \right) \widehat{D}\vec{D}_n \right) - \left(\frac{K_3}{|\vec{D}\vec{W}|^2} \right) \widehat{D}\vec{W} - \left(\frac{K_4}{|\vec{D}\vec{R}|^2} \right) \widehat{D}\vec{R}$$

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

Fig. 4. Flock model (schematic not drawn to scale). Key: gain parameters $K_{1 \rightarrow 4}$; 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 \rightarrow 4) and resultant velocity \vec{d} (where \hat{a} is the unit vector of \vec{a}).

which produces repulsion from the robot is proposed to model the aversive response of the ducklets to the robot. Note that all these forces are scaled according to the inverse square of distance. Each ducklet moves according to the resultant of the forces acting upon it, subject to a simulated inertia that smoothes acceleration, and limited by a top speed chosen to approximately match that of the real ducks. 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 Experiments

Experiments with the simulator guided the development of two novel flock control algorithms which are closely related to the flock model described above. Due to space limitations, only the most recent and successful of these ('Method




Fig. 5. Robot controller(schematic not drawn to scale). Key: gain parameters $K_{1,2}$; flock centre F; Robot position R; Flock goal position G; algorithm terms (1 \rightarrow 3) and resultant \vec{r} (where \hat{a} is the unit vector of \vec{a})

2') is presented here. The original method is described along with our first results in [14].

The distance $|GF|$ in Figure 5 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 \vec{r} is given by the function shown in Figure 5. 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. Note the simplicity of the algorithm and that it is expressed in similar terms to the flock model.

4.1 Performance in simulation

The algorithm is first tested in simulation. A point on the arena boundary is chosen as the flock goal, twelve 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 centre 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.

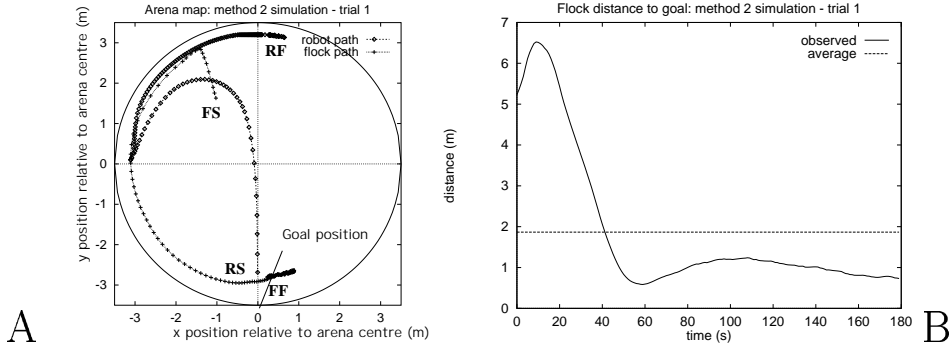


Fig. 6. Simulation results: (A) paths in arena [Key: RS = robot start, RF = robot finish, FS = flock start, FF = flock finish] and (B) distance to goal over time.

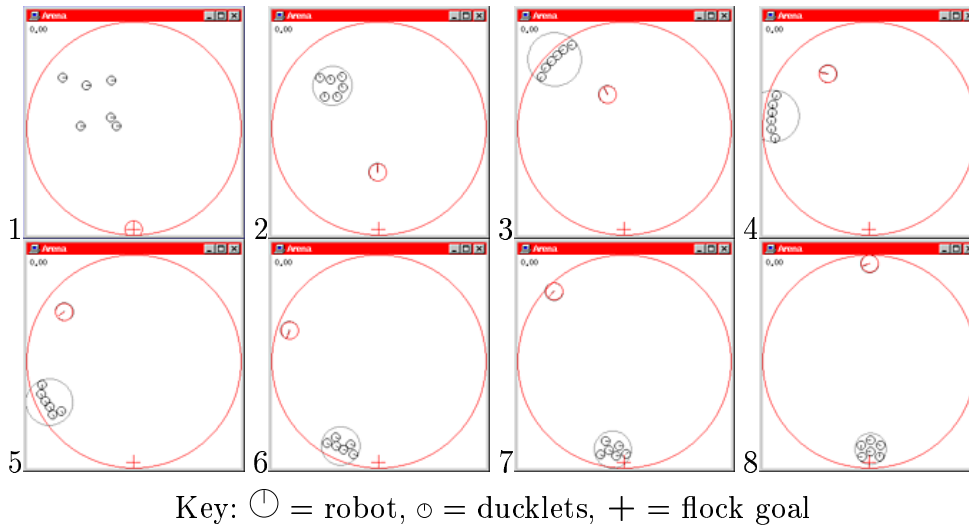


Fig. 7. Sequence of images from the simulator during a trial.

The results show that this controller successfully performs the required task. Figure 6 (A) shows a representative plot of the simulated robot and flock paths around the arena, while Figure 7 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 trial's success for comparison 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, with a standard deviation of 0.16m.

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

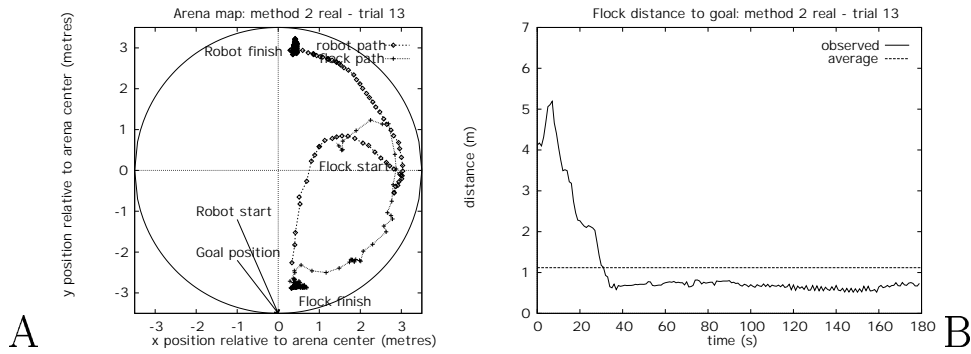


Fig. 8. Real-world results: (A) paths in arena and (B) distance to goal over time.




Fig. 9. Sequence of images from the overhead camera during an experiment. The goal position is at the bottom of the picture.

and positioned near the goal, a flock of twelve ducks is introduced into the arena. For three minutes the ducks were left to become accommodated to the arena. Their positions were not recorded during this time, but they were typically observed to settle into a stationary, loosely-aggregated group with no common orientation. The settled position of the flock varied apparently at random between trials. After three minutes the robot is activated. The positions of the robot and flock centre 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 seven times with the same flock, with the robot at a slightly different position near the flock goal in each trial. All the ducks were the same age and had been raised under similar conditions (described completely in [4]).

Figure 7 (A) shows a representative plot of the real robot and flock paths around the arena, while Figure 9 shows 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 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 1.12m. The average score over all seven trials was 1.67m with a standard deviation of 0.28m.

The larger average score in the real world compared to simulation was largely 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 method varies (slightly) with this setting, there is scope to devise further algorithms that may be more robust with respect to the inevitable variation between flocks.

5 Scope, limitations and extensions

The robot task we have examined was deliberately restricted to allow us the maximum chance of success in the three-year life of the project, while still offering an interesting demonstration of robot/animal interaction. Having established that we can reliably perform this minimal task, there are many possible variations and extensions that could be examined.

5.1 *Flock-splitting and unused information*

Our controllers do not use the flock radius obtained by the vision system. This is a potentially useful piece of information that could be used to make the robot more responsive to flock state. The main reason for an increase in flock size would be when the flock splits into two or more sub-flocks, or when an individual breaks from the main flock. Flock animals strive to maintain proximity, so a separation can be considered a stressful event that should be avoided. Monitoring the size of the flock could allow the robot to back away when a split happened, allowing the animals to re-form a single flock.

Conversely, a split may be desirable, for instance if the task were to isolate individuals for inspection or veterinary attention. A sudden increase in flock size could indicate to a more sophisticated vision system that a single flock should be re-assessed as multiple sub-flocks or individuals.

The speed of the flock, though easily obtained from the existing tracking data, is not exploited here. A more sophisticated flock control strategy could be based on the relative *velocities* of the robot and flock, rather than the *positions* used up to now. This could perhaps allow more subtle control to further reduce animal stress or to improve absolute performance.

5.2 *Free-space goals*

The robot currently holds a flock close to a point against a wall. If the goal position is moved into free space (ie. there is no wall, or the wall is too far away to significantly effect the flock or robot) then the system's behaviour changes. Testing this scenario in the simulation the flock is observed to orbit the goal position, with the robot orbiting the goal at some greater distance. Depending on the relative speeds of the flock and robot as they first approach the goal, the orbit may be simple and elliptical or it may be complex and apparently chaotic. However, provided the robot can move faster than the flock, the flock is always contained within some finite area centred at the goal. Some further mechanism is required to ensure that the flock reaches a goal in free space and *stops* there.

5.3 *Corners*

To control the movement of a flock in a more realistic environment, such as a poultry house or farmyard, the robot would have to deal with corners. Corners are often a problem for potential-field controllers as they are a common source of local minima. The 'cornered animal' is similarly stuck in a local minima if its situation is undesirable but any movement makes the situation worse.

Corners may require substantial extension and/or modification of the controllers so far developed. Two simple observations may guide the development of a corner-capable system: (1) to ensure an animal or flock moves out of a corner, the robot must move into the corner. Assuming the robot does not physically block their escape, they will eventually move; (2) if the robot approaches from one side rather than the middle of the corner, then the animal or flock is likely to leave by the other, open side.

5.4 *On-board sensing*

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. Again, this arrangement

was chosen at the start of the project for pragmatic reasons. 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 dietic sensor modes without having to physically engineer them. We have devised similar flock control algorithms, based on the earlier Method 1 controller, which use only local sensing (ie. range data and/or vision) in to achieve the same task. We intend to similarly adapt the more successful Method 2 controller described in this paper to these on-board sensing schemes. The results of this work will be reported elsewhere.

5.4.1 *Adaptation for improved performance*

This work has approached the issue of variation between flocks by designing a robust, general flock-control method. While it is successful in a majority of trials, its success and efficiency varies between flocks and over time with the same flock. While it was one of our goals that the controller should not require manual optimization to work with any specific animal or group, the ability to self-optimize or *adapt* during run-time could be a very useful extension. Enhancing a general strategy with adaptation might allow a robot to immediately interact with a novel target flock in approximately the right way, with its performance improving over time as its experience increases.

6 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.

Acknowledgements

Thanks to Steve Crook, Paul Griffith, John Lowe and Robert Smith for their work on the robot, and to the SIRS'98 conference organisers and reviewers. The Robot Sheepdog Project was funded by the UK's Biotechnology and Biological Sciences Research Council and Silsoe Research Institute.

References

- [1] R. Boyle, V. Hlavac, and M. Sonka. *Image Processing, Analysis and Machine Vision*. Chapman and Hall, 1993.
- [2] S. Cameron and P. Probert. *Advanced Guided Vehicles: Aspects of the Oxford AGV project*. World Scientific, 1994.
- [3] W.D. Hamilton. Geometry for the selfish herd. *Journal of Theoretical Biology*, 31:295–311, 1971.
- [4] Jane V. Henderson. *Adaptive Responses of Herding Animals to a Robot Vehicle*. PhD thesis, University of Bristol, expected 1999.
- [5] M.J. Mataric. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1):51–80, 1996.
- [6] B. L. Partridge. The structure and function of fish schools. *Scientific American*, pages 114–123, June 1982. June 1982 issue.
- [7] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioural model. *Computer Graphics*, 21(4):25–34, jul 1987.
- [8] N. Sumpter, R.D. Boyle, and R.D. Tillet. Modelling collective animal behaviour using extended point distribution models. In *Proceedings of the British Machine Vision Conference*, 1997.
- [9] N. Sumpter and A. Bulpitt. Learning spatio-temporal patterns for predicting object behaviour. Technical report, University of Leeds, School of Computer Studies, The University of Leeds, UK, 1998.
- [10] N. Sumpter, A. Bulpitt, R. Vaughan, R. Tillet, and R. Boyle. Learning models of animal behaviour for a robotic sheepdog. In *Proceedings of the IAPR Workshop on Machine Vision Applications*, 1998.
- [11] R. Vaughan, J. Henderson, and N. Sumpter. Introducing the robot sheepdog project. In F. Juste, G. Andreu, J.M. Valiente, and J.V. Benlloch, editors, *Proceedings of the International Workshop on Robotics and Automated Machinery for BioProductions*, September 1997.

- [12] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Experiments in automatic flock control. In *Proceedings of the Sixth Symposium on Intelligent Robotic Systems (SIRS'98)*, 1998.
- [13] 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, 1998.
- [14] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic animal control. In *From Animals to Animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behaviour*. MIT Press, 1998.
- [15] K. Warburton and J. Lazarus. Tendency-distance models of social cohesion in animal groups. *Journal Of Theoretical Biology*, 150:473–488, 1991.


Biographies

Richard Vaughan



Richard Vaughan expects to gain a DPhil. from the University of Oxford in 1999 for his work on the Robot Sheepdog Project. He is interested in how animals do what they do, and in their lessons for robustness and autonomy in robots. He is currently a Research Associate at the University of Southern California's Robotics Research Laboratory.

Neil Sumpter




Neil Sumpter expects to receive his PhD. from the University of Leeds in 1999 for his work on recognition, tracking and modeling of groups of animals as part of the Robot Sheepdog Project.

Jane Henderson



Jane Henderson graduated from the University of Edinburgh with an MSc. in Applied Animal Behaviour and Animal Welfare in 1995. She expects to gain her PhD. in Animal Behaviour from the University of Bristol in 1999. Her thesis on flock behaviour and enemy assessment forms part of the Robot Sheepdog Project.

Andy Frost



Andy Frost is a Mechanical Engineer and is Leader of the Livestock Engineering Group at Silsoe Research Institute, specialising in the development of systems to interact with animals.

Stephen Cameron



Stephen Cameron is a Reader in Computing Science at Oxford University Computing Laboratory, and Fellow of Keble College. He obtained his PhD in Artificial Intelligence at Edinburgh University, working on the geometric modeling of robots and on collision detection. His research spans a broad area of spatial reasoning, including its application to robot planning and control.

He is a member of the AISB, the ACM, the IEEE Robotics and Automation society, and the Geometric Modelling Society.