Introducing the Robot Sheepdog Project

Richard Vaughan[1,2], Jane Henderson[1,3], Neil Sumpter[1,4]

[1]BioEngineering Division, Silsoe Research Institute, Silsoe, Bedford MK45 4HS, UK.
[2]Computing Laboratory, University of Oxford (vaughan@comlab.ox.ac.uk);
[3]Dept. of Clinical Veterinary Science, University of Bristol (jane.henderson@bbsrc.ac.uk);
[4]School of Computer Studies, University of Leeds (neils@scs.leeds.ac.uk).

Abstract

This paper introduces the Robot Sheepdog Project; the first experiment in Animal-Interactive Robotics. The goal of the project is to produce an autonomous robot that controls the movement of a flock of ducks. The robot system and its vehicle is described, and its control and machine vision modules are outlined. Three complementary studies of duck flocking behaviour are described. One of these flocking models inspired a prototype flock-control method which is presented here for the first time.

Key words: robotics, animals, interaction, flock, vision, behaviour, sheepdog.

1 Introduction - Fresh AIR

Robots have already found a place in animal husbandry systems through automatic milking systems and robotic sheep shearing. It is hoped that robots could play an increasing role in agricultural systems, replacing humans in hazardous, tedious or unpleasant tasks, or where there are benefits in efficiency, effectiveness and animal welfare. Thus, robots may be especially useful for inspection and handling of animals, and waste management.

The task of designing robots to interact with animals raises issues not encountered in other industries. First, the welfare of animals must be considered; robots must be designed to cause no unecessary stress to the animals they work amongst. Secondly, unlike the objects of typical robot applications, animals are autonomous agents and will exhibit behaviour. This is both a problem and an opportunity. For example most animals can move themselves around; they may not stay where you left them, but on the other hand they may move to where you want them. Thus animals need not be handled directly, but can (and perhaps should) have their *behaviour* manipulated towards achieving some goal. We believe that such Animal-Interactive Robotic (AIR) systems are an important area for research in BioRobotics, and our project was conceived to test these issues.

2 The Robot Sheepdog Project

We present the Robot Sheepdog Project; an experiment in Animal-Interactive Robotics. We aim to demonstrate a robot system that will enter an arena, gather a flock of ducks and manœuvre them safely to a predetermined goal position. Ducks flock similarly to sheep, but their small size and low speed simplify our experimental requirements. The problem remains essentially the same. A pilot study showed that a human operator could achieve this task (with some difficulty), using a radio controlled model car.

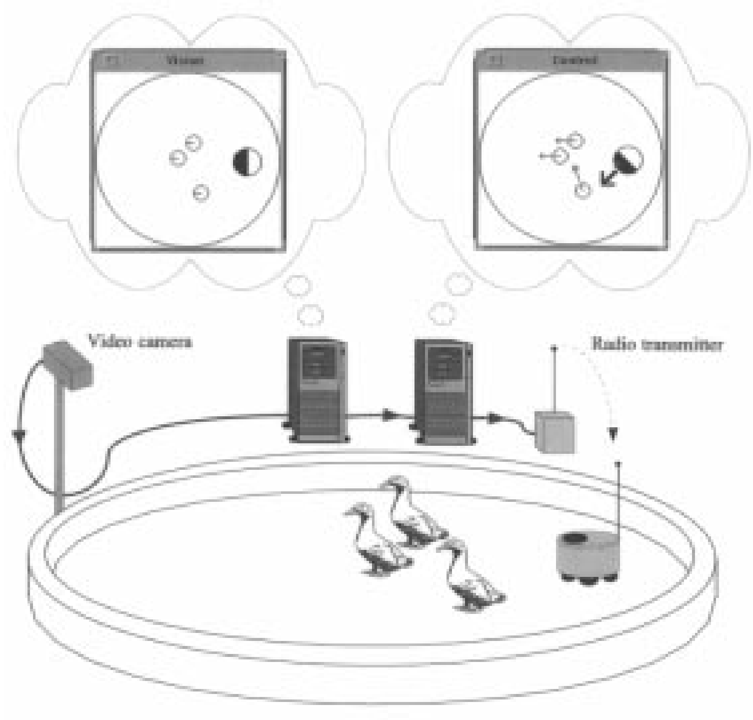

Figure 1: Robot Sheepdog System Overview

We identified these elements to be required for a robot sheepdog system:
- a vehicle to interact appropriately with the ducks (see section 3);
- a means of determining the current positions of the robot and ducks (see section 4);
- knowledge of the ducks' responses to the vehicle (see section 5).
- an algorithm to control the vehicle and effectively herd the ducks (see section 6).

This paper presents an outline of the project so far.

3 Rover the Robot Sheepdog

The robot sheepdog system comprises a custom-built vehicle, two workstations and a camera (Figure 1). The vehicle and the ducks operate in an arena of about 7m in diameter, in view of the overhead video camera. One workstation analyses the images from the camera to find the position and orientation of the robot and flock (see section 4). A high-level 'strategic' control program runs on another workstation, generating a vehicle speed and turn-rate, guiding the vehicle to control the flock (see section 6). The design of this controller incorporates information about the responses of the ducks to movements by the robot (see section 5). The robot has its own low-level PID controller which adjusts the power to the wheels in order to achieve the required behaviour. The vehicle moves, the ducks move in response, and the new positions are observed by the camera. By distributing the system in this way we can achieve high processing speed using only low-bandwidth communications between modules.

The vehicle is required to work in a duck's environment: outdoors, on short grass, and in real time. Thus our robot has a top speed of around $4ms^{-1}$ and acceleration of $\pm 2ms^{-2}$, which is about twice as fast as our 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*.

The system is assembled and currently *Rover* finds its way about the arena and can follow a

Figure 2: Finished *Rover the Robot Sheepdog*, with cover removed.

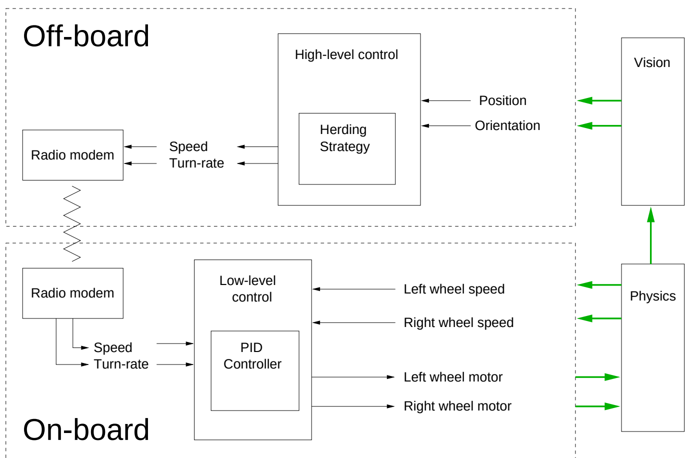

Figure 3: *Rover*'s control schematic

human around, standing off to a specified distance. Work has begun on developing flock control strategies, first in simulations incorporating our duck models (section 6), then tested with the real ducks. See [9] for details of this work.

4 Machine Vision for animal location and vehicle guidance

The flock is located and the robot controlled using machine vision techniques. The large area viewed from the overhead camera means that is it very difficult to resolve individual birds. Thus we identify and track the *flock* of ducks as one entity, and pass the center and average radius of this object to the high-level robot controller, rather than the location of individual ducks.

In order to quickly develop a working robot, a simple method of combining standard image processing and feature location strategies is used. A reference image of the scene background is learned, and this is subtracted from the incoming images from the camera in order to distinguish movement in the scene. Hysteresis and region growing are then applied to determine the most significant regions of activity within the image (see Figure 4). Such techniques are described in standard textbooks (e.g. [1]).

Figure 4: (a) Typical image taken from video sequence, and (b) Segmented robot and flock regions

Additionally, a marker on top of the vehicle - out of view of the animals, but observed by the camera - can be used to measure vehicle orientation. We paint the circular lid of the robot half black, half white, and detect the angle of the high-contrast line in the middle. We can detect *Rover*'s orientation to approximately ± 4 degrees.

Successful robotic control is often reliant on a high information rate from its sensors. Currently the image analysis can generate the required positional data at 15hz, to an accuracy of approximately ± 20cm.

5 Flock behaviour and visual perception

There are very few studies of the behaviour of animals in the presence of machines or robots. Although clearly robots do not form part of an animal's evolutionary precepts, studies have shown that cows readily adapt to robotic milking systems within a few days [6], and mechanised broiler catchers have been shown to evoke lower stress type responses than manual collection [3]. Animals may indeed find it easier to adapt to more predictable automated systems than to (often inconsistent) stockman behaviour, and automated sytems have potential to allow animals a greater degree of control over some aspects of husbandry.

5.1 Ethological data

Behaviour experiments have been divided into two areas. The first set of experiments investigates the responses of ducks when confronted with a potential threat (e.g. human, robot), and provides information regarding flock dynamics (e.g. parameters such as flight distance, and escape path), individual variation, and flock variation. Recent experiments have found a highly significant tendency (Kendall's Concordance, $p < 0.005$) for individuals in small flocks of ducks to maintain a relative position within the flock; and that ducks habituate rapidly to a novel moving stimulus. This sort of information may prove important to the design of an effective robot flock control system.

A second set of experiments investigates how ducks perceive a stimulus according to its visual appearance; investigating the relative importance of certain features of a mammalian predator (e.g. posture, eyes and ears), and how such features make a stimulus appear more or less threatening to a prey animal. Results from these experiments will provide a basis from which to design *Rover*'s appearance according to the degree of control (or threat) needed over the animals at a particular time, or in a particular situation. See [5] for details of this work.

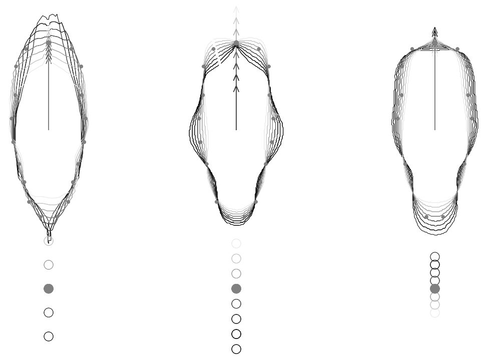

Figure 5: Examples of variations in the shape model. The arrow represents the flock velocity, and the circle represents the robot position

5.2 Building shape models of animal flocks from vision data

We also apply machine vision to help us understand the dynamics of the flock as it moves to avoid the robot.

We consider how the robot location can affect the shape of the flock, which may lead to a simple model of behaviour. For example, as the robot approaches the animals they take fright, changing the shape of the flock. Thus if we can accurately capture the relationship between flock shape and robot position in a sequence of images, we can create a model that describes some behavioural aspects of the animals.

A statistical method for automatically creating such a model from a sequence of images is described by one of the authors in [8], based upon Point Distribution Models [2]. The shape of the duck flock is expressed in terms of its outline boundary, and additional parameters can be included such as the flock velocity which are also dependent on the robot position. Examples of the variations in such a shape model can be seen in Figure 5.

The model can be applied to the problem of tracking, where the processing speed can be increased due to a knowledge of the typical shape of the desired features. Additionally, if the model can accurately predict the motion of the flock, it may be used to further extend the model of flocking which determines the herding strategy of the robot.

5.3 A minimal model flock

Models of flocking behaviour have been suggested by ethologists [4] and constructed by computer animators [7] to produce spectacularly realistic animations of flocking birds, bats (Batman Returns), wildebeast (The Lion King) and dinosaurs (Jurassic Park). The flock entity we perceive can be produced by the *mass uncoordinated action of individual flock members*. Each animal in the flock performs its own (possibly very simple) behaviour in response to its local environment, with the global effect of producing and maintaining the flock.

We are interested in the movement of ducks in our arena in response to a perceived threat from the Robot Sheepdog. Realistic-looking movement can be observed in simulated ducks that move in response to the following simultaneous 'forces' or behavioural 'drives':
1) Attraction to other ducks by an amount proportional to their mutual distance; 2) Repulsion from other ducks by an amount proportional to the inverse square of their mutual distance; 3) Similar repulsion from obstacles; 4) Similar repulsion from the threat.

These forces are balanced so that the model duck (or *ducklet*) moves towards other ducks (duck attraction), but maintains a small distance from them (duck repulsion). As a threat

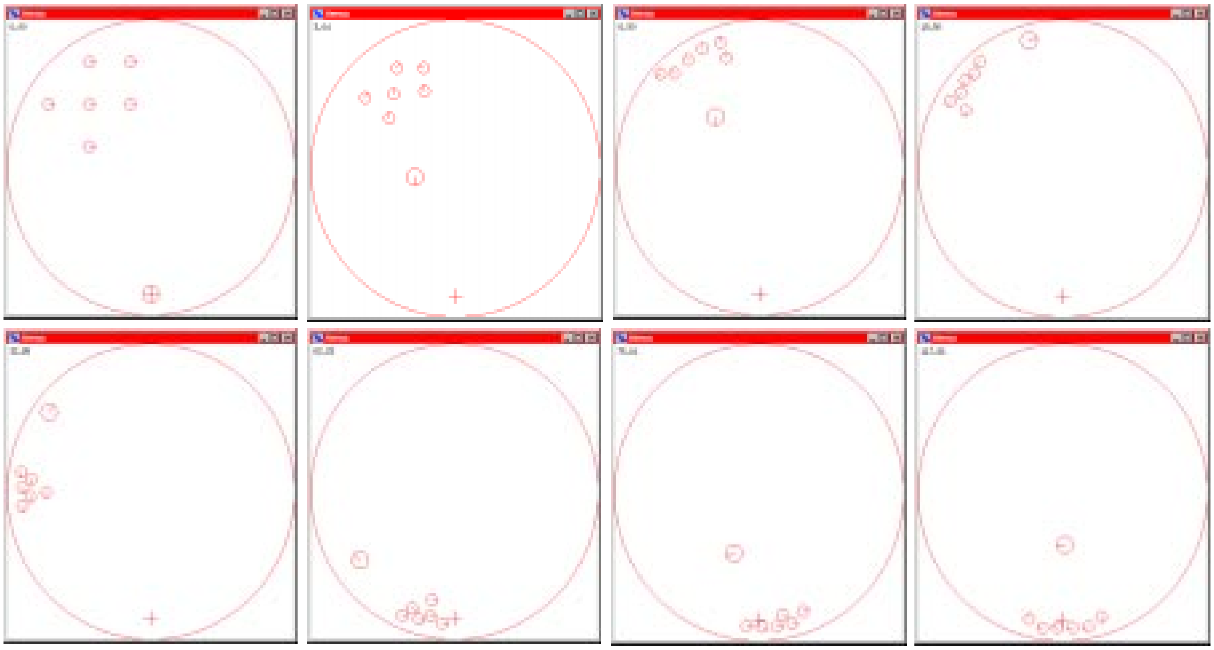

Figure 6: Sequence of DuckSim screenshots showing the simulated *Rover* herding the ducklets. The sequence reads from left to right, top to bottom. The goal position is indicated by the '+' symbol.

approaches it will move away (threat repulsion) whilst avoiding any other obstacles (obstacle repulsion). This method is essentially a potential field algorithm.

Flock characteristics such as inter-animal spacing, speed of individuals, speed of flock movement, etc., are determined by the absolute and relative magnitudes of the four forces. We aim to optimise these parameters to best match the behaviour of our real ducks.

The simulator is used as a tool in the design and testing of candidate flock control strategies. This provides all the advantages of a typical engineering simulation, plus it minimizes the use of the real ducks. Later work will compare this general minimal flock model to the vision-based statistical model described above.

6 Controlling the Flock: a candidate strategy

A method for controlling the flock was devised while experimenting with the duck simulator. The technique exploits the nature of the flocking model, and is built from the same components of attractive and repulsive fields. This commonality of mechanism between flock and flock-controller is intuitively attractive, and enables a very simple technique to perform very well, at least in the simulator. Figure 6 shows the simulation running. Recall that the task is to fetch a group of ducks from the arena back to a goal position. The robot starts out at the goal and moves towards the ducklets (top left). They form into a flock as the robot approaches, and move directly away from it until they reach the boundary. The robot moves behind them with respect to the goal and then pushes them towards it. As the ducklets approach the goal the robot moves around to stop them overshooting, then stays at a constant stand-off distance to hold the flock at the goal (bottom right).

6.1 How it works

The robot's action at each time step is determined by a potential field algorithm in the same way as the simulated ducks:
1) Attraction to ducks by an amount proportional to their mutual distance; 2) Replusion from ducks by an amount proportional to the inverse square of their mutual distance; 3) Similar repulsion from obstacles; 4) Repulsion from the goal position by a constant amount.

These forces are balanced so that *Rover* moves towards the flock (duck attraction), but maintains a set distance from them (duck repulsion). The sum of these two forces produces an 'orbit' of minimum potential around the flock, with the potential landscape resembling a Mexican hat. The effect of the goal-repulsion is to tilt the whole landscape - lifting the hat at one side - which produces a single point of minimum potential *behind the flock with respect to the goal*. If the balance of forces is such that this point lies within the flight distance of the ducks then they will move directly away from the *Rover* and thus towards the goal point. A further force acts to keep the robot from colliding with obstacles and the boundary walls. The potential field thus generated is used to determine the ideal position for *Rover* to be in. The control system drives *Rover*'s wheels to attempt to attain that position.

Initial simulation trials show very robust performance, with low sensitivity to parameter settings and sensor noise. We believe the efficacy of the method is due to its close relationship with the mechanism of the model flock. If the model captures the essentials of flocking behaviour, as we hope it does, the technique has a good chance of working with real animals. Experiments in the next few weeks will test the method on real duck flocks.

7 Acknowledgements

Thanks to our supervisors: Christopher Wathes, Stephen Cameron, Roger Boyle, Robin Tillett, Andy Frost, Christine Nicol, Jeff Lines and Nick Tillett.

References

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

[2] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Training models of shape from sets of examples. In *Proceedings British Machine Vision Conference*, pages 9–18, 1992.

[3] I. Duncan, G. Slee, P. Kettlewell, P. Berry, and A. Carlisle. Comparison of the stressfulness of harvesting broiler chickens be machine and by hand. *B. Poultry Science*, 27:109–114, 1987.

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

[5] J. V. Henderson. Adaptive responses of animals to mobile robots. Technical report, University of Bristol, 1997.

[6] N. B. Prescott. *Dairy cow behaviour and automatic milking*. PhD thesis, University of Bristol, 1995.

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

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

[9] R. Vaughan. Experiments in animal interactive robotics. Technical report, Computing Laboratory, University of Oxford, 1997.