MVA '98 IAPR Workshop on Machine Vision Applications, Nov. 17-19, 1998, Makuhari, Chiba, Japan

15—3

Learning Models of Animal Behaviour for a Robotic Sheepdog

N. Sumpter[1,2], A.J. Bulpitt[1], R. Vaughan[2,3], R.D. Tillett[2] and R.D. Boyle[1]

[1]School of Computer Studies, The University of Leeds, Leeds, UK
[2]Silsoe Research Institute, Wrest Park, Silsoe, Beds., UK
[3]Oxford University Computing Laboratoty, Parks Road, Oxford, UK
e-mail: `neils@scs.leeds.ac.uk`, `andyb@scs.leeds.ac.uk`

Abstract

To further the use of machine vision in animal-related tasks such as automated monitoring, an understanding of the behaviour of the animals in their environment is required. This paper describes an application, the Robotic Sheepdog, which exploits animal behaviour to achieve its goal. We present a method of automatically extracting a model of animal behaviour that is deemed more appropriate than an alternate rule-based solution, and describe how this can be used to determine a likelihood of future events.

1 Introduction

The application of machine vision techniques to the domain of animals is relatively unexplored compared with other areas of computer vision research, despite its obvious potential. Applications such as Schofield's [1] use of a vision system to determine automatically a pig's weight by the surface area of the pig observed in an image could prove very useful to the agricultural industry.

One reason for a lack of animal related applications is an insufficient understanding of the animals behaviour within their environment. Simple rule-based systems of animal behaviour can be used to give visually appealing results (for example, animations of animal flocking [2]), but these do not necessarily portray true behaviour.

In this paper, we describe how a model of animal behaviour can be learnt automatically from image sequences, and applied to predict future behaviour from a recent history of animal motion.

The work forms part of the Robotic Sheepdog Project [3]; an investigation into animal interactive robotics. The aim of the project is to demonstrate an autonomous robot system that can successfully manipulate a group of animals to some pre-determined goal, by exploiting the animals' adaptive behaviour. As a robotic task this differs significantly from previous work combining robots and animals, where animal behaviour is deliberately minimised by physical restraint [4]. The task of the sheepdog was chosen due to 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.

In order to successfully exploit the behaviour inherent to the animals, a model of the likely reaction of the flock to the robot vehicle must be constructed. Such a model can be built using a rule-based solution to provide a control strategy for the robot [5]. However, whilst this provides a simulation of animal flocking that is visually similar to the real animal behaviours, it is argued that a model learned automatically from observations (in terms of image sequences), will provide additional information of the real environment that a rule-based approach would not be capable of encapsulating.

2 Method

In describing the behaviours of the animals, we consider modelling the location and velocity of both the flock and the robot, together with the flock shape. The flock is modelled as a whole, not as individual birds. This is due to the poor resolution obtained when fitting the whole environment of the animals within the image frame. The flock and robot move within a circular arena, eight metres in diameter (see Figure 1).

Image frames are subtracted from a pre-learned background image, significant regions extracted and

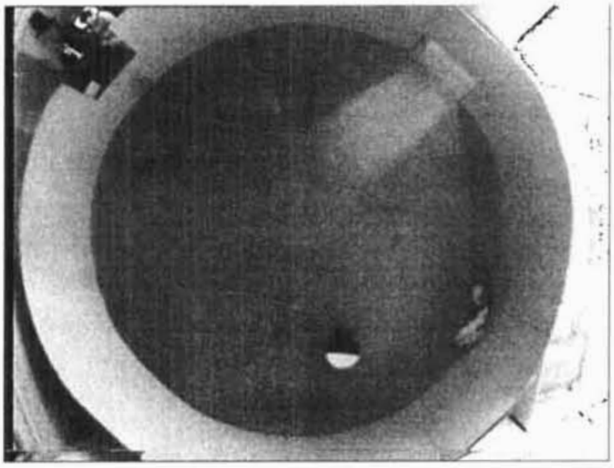

Figure 1. A typical image of the arena.

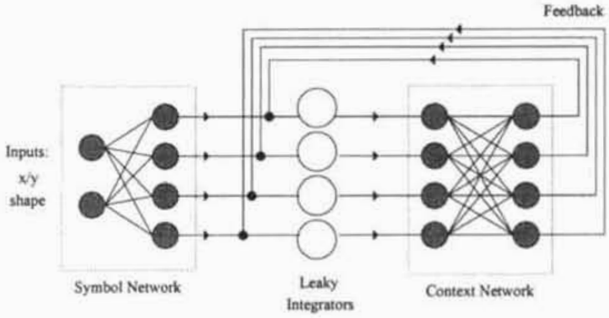

Figure 2. The approach used, represented as a network architecture

then subjected to morphological smoothing. This provides a segmentation scheme for the flock as a group. The robot can be found by a high-contrast black and white motif placed on the top of the vehicle; the design also enables us to determine the robot's orientation.

The shape of the flock is also of interest, because it represents behavioural traits of the animals; for example, a long elliptical flock shape indicates panic as the animals flee from the robot predator. The shape can be modelled by using the outline of the segmented flock region as the basis for a Point Distribution Model (PDM) [6] to reduce the dimensionality of the shape data.

By combining the locations of the flock and robot with the appropriate principal shape parameters of the PDM, a scene vector $\mathbf{x}$ can be constructed for each frame in an image sequence.

The model of animal behaviour is a representation of the spatio-temporal patterns of the animals within their environment, represented for time t as the conditional probability $p(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{x}_{t-1}, \ldots, \mathbf{x}_{t-n})$, allowing the implicit generation of plausible future motions and appearance changes given recent observations.

To estimate this conditional probability density function (pdf), a state-based approach is used where a temporal sequence is considered in terms of the *symbols* observed (the feature vectors) and the *context* in which they appear. Thus the pdf becomes dependent on the current feature vector $\mathbf{x}_t$ and a representation of the contextual history $\mathbf{H}_t$. To achieve this, a neural network architecture is used as in Figure 2. Two competitive learning networks [7] are connected by a layer of leaky integrators [8].

The symbol network is used to perform vector quantisation on the space of the input scene vectors, $\mathbf{x}_t$, effecitively clustering the features into a set of M symbol states.

The outputs of the symbol network are connected to a layer of leaky integrators. Each integrator is essentially a decay function: the winning symbol output (according to the competitive learning algorithm) causes the associated integrator value to rise; when the symbol output no longer wins, the integrator tails off slowly over time. In this way, by examining at time t the relative values on all the integrators, $\mathbf{H}_t$, a notion of the order in which previous symbols were observed is obtained.

The context network performs vector quantisation upon the space of the leaky integrator values, but modified so that the next symbol output is associated with the previous integrator state. The effect of this is that an observed leaky integrator representation of history $\mathbf{H}_t$ produces a high response for the next symbol in the sequence, $\mathbf{x}_{t+1}$. The inclusion of a feedback loop, similar to that of [9], makes this association implicit, and enforces that the estimated probability of state at a given time is dependent on both the currently observed symbol as well as the recently observed history.

This approach is similar to that of Johnson and Hogg [10], who use a comparable architecture to model human trajectories in order to monitor typical and atypical events. Our approach improves on this with the inclusion of the feedback mechanism to enable implicit prediction of future trajectories without the need for an extra learning phase. Full details of the method used to train such a model are given in [11].

3 Evaluation and Results

The network is trained on a set of 20 sequences that represent typical behaviours of the animals. Each sequence consists of between 400 and 1200 frames, with

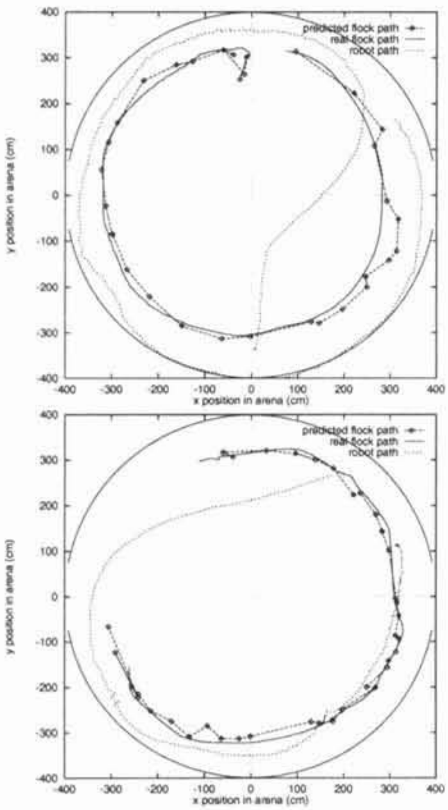

Figure 3. Results of prediction in comparison to actual flock path, (top) for a sequence used to train the model and (bottom) for an unseen path

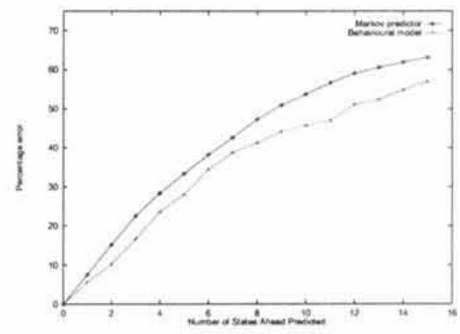

Figure 4. State mis-classifications between prediction and actual sequence

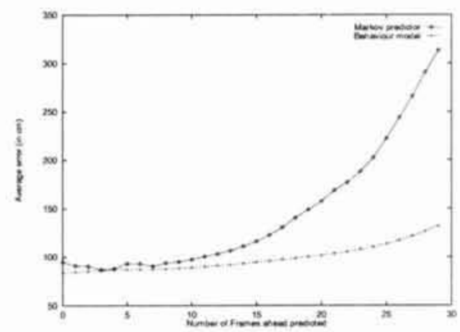

Figure 5. Average distance error between predicted and actual paths

each sequence beginning at different positions within the arena.

One of the principal aims of the approach taken is to be able to present partial information to the model in order that the most likely corresponding missing data is obtained, for instance presenting a known robot location and generating the most plausible location for the animals. Figure 3 shows typical qualitative results of prediction, using the trained model. The path of the robot from a sequence is presented to the network, and the corresponding predicted path of the flock is shown. For both training sequences and unseen paths, it is observed that the predicted behaviour closely represents the original path.

A more suitable method of evaluation can be achieved by comparing the predicted sequence of symbol states (obtained by the maximum likelihood valued context output) with the actual symbol states that a known sequence passed through. Figure 4 shows the percentage of incorrectly matched states against the number of states ahead predicted. For comparison the results for a simple Markov chain are also presented, this being modelled as the one-step transitional probabilities between symbol states, again with the maximum value probability being chosen. It is observed that the behavioural model produces better results than the Markov process indicating that the historical evidence is producing a better prediction decision.

In the above evaluation, however, it is conceivable that when a mis-classification between predicted and actual state occurs, the (wrongly) predicted state lies close to the original state in the real feature space. Thus Figure 5 shows the average distance error between predictions and real sequences. The minimum error distance of approximately 50cm occurs due to quantisation and calibration effects. It is noted that the model performs very well in comparison to the Markov chain, which diverges greatly from the actual location as time increases. The error distance increases also over time, but only by around 50cm over 30 frames (approximately 3 seconds), which is the diameter of the robot.

4 Future Work and Conclusions

The general framework of the model to predict the next state given current information (in terms of the observed symbol and temporal context) lends itself well to Isard and Blake's CONDENSATION [12] algorithm, and future work concerns the incorporation of the model into this tracking paradigm. Heap and Hogg [13] describe how a Markov chain between areas of shape-space can be combined with CONDENSATION to good effect, and since the results presented here show a distinct improvement over the kind of Markov predictor used, it is reasonable to expect the presented model will provide a stronger tracking mechanism.

The approach taken also has potential for providing an automated control strategy for the robot itself, since the model inherently predicts not only the path and shape of the flock, but also of the robot. A goal location for the flock can be chosen, and a path of maxmimum likelihood from the current position to the goal found. The corresponding predicted robot path represents the best path for the robot to follow if the ducks are to be successfully herded to their goal.

In conclusion, this paper presents a machine vision application where animal behaviours are learned automatically from image sequences. A model of such behaviours proves essential in this unique robot scenario - using an autonomous vehicle to herd animals to a goal. The method for learning behaviours is novel, and provides an approach which is more appropriate than rule-based alternatives. Suitable extensions to the work are discussed, illustrating the applicability of the model to the tasks of tracking and control.

5 Acknowledgements

The authors would like to thank both EPSRC and BBSRC in the UK for the funding of this research.

References

[1] C.P. Schofield. Image analysis for non-intrusive weight and activity monitoring of pigs. *Proc. 4th International Symposium on Livestock*, pages 503–510, 1993.

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

[3] R. Vaughan, J. Henderson, and N. Sumpter. Introducing the robot sheepdog project. In *Proc. Int. Workshop on Robotics and Automated Machinery for Bio-Productions*, 1997.

[4] A. Frost, M. Street, and R. Hall. The development of a pneumatic robot for attaching a milking machine to a cow. *Mechatronics*, 3(3):409–418, 1993.

[5] R. Vaughan, N. Sumpter, A. Frost, and S. Cameron. Robot sheepdog project achieves automatic flock control. In *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*, 1998. To appear.

[6] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Training models of shape from sets of examples. In *Proc. 3rd British Machine Vision Conference*, pages 9–18, 1992.

[7] D.E. Rumelhart and D. Zipser. Feature discovery by competitive learning. *Cognitive Science*, pages 75–112, 1985.

[8] M. Reiss and J.G. Taylor. Storing temporal sequences. *Neural Networks*, 4:773–787, 1991.

[9] A.J. Bulpitt. *A Multiple Adaptive Resonance Theory Architecture Applied to Motion Recognition Tasks*. PhD thesis, Dept of Electronics, University of York, 1994.

[10] N. Johnson and D.C. Hogg. Learning the distribution of object trajectories for event recognition. *Image and Vision Computing*, 14(8):609–615, 1996.

[11] N. Sumpter and A.J. Bulpitt. Learning spatio-temporal patterns for predicting object behaviour. In *British Machine Vision Conference*, volume 1, September 1998. to appear.

[12] M. Isard and A. Blake. Contour tracking by stochastic propagation of conditional density. In *European Conf. of Computer Vision*, pages 343–356, 1996.

[13] T. Heap and D.C. Hogg. Wormholes in shape space: Tracking through discontinuous changes in shape. In *Proc. Sixth International Conference on Computer Vision*, pages 344–349, Bombay, India, 1998.