
Online Collective Animal Movement Activity Recognition

Kehinde Owoeye

Department of Computer Science
University College London
London, WC1E 6EA
ucabowo@ucl.ac.uk

Stephen Hailes

Department of Computer Science
University College London
London, WC1E 6EA
s.hailes@ucl.ac.uk

Abstract

Learning the activities of animals is important for the purpose of monitoring their welfare vis a vis their behaviour with respect to their environment and conspecifics. While previous works have largely focused on activity recognition in a single animal, little or no work has been done in learning the collective behaviour of animals. In this work, we address the problem of recognising the collective movement activities of a group of sheep in a flock. We present a discriminative framework that learns to track the positions and velocities of all the animals in the flock in an online manner whilst estimating their collective activity. We investigate the performance of two simple deep network architectures and show that we can learn the collective activities with good accuracy even when the distribution of the activities is skewed.

1 Introduction

Recognising the collective movement activities of animals is important for several real world applications such as monitoring their welfare vis a vis their behaviour with respect to their environment and conspecifics, predicting the onset of an epidemic or attacks from predators most especially for animals who live in cooperative societies. While several works have been carried out in recognising collective activities with respect to human interactions [1, 2], previous works in the animal behaviour community however have either tried to recognize collective behaviour from different species of animal where a model is fed the entire collective behaviour input and asked to identify the specie generating such behaviour [3] or learn the activity of just a single animal [4, 5]. There are however problems with these approaches. Learning collective activity in an offline manner doesn't scale into real world applications where the collective activity is required real time for example in monitoring poaching activities. In addition, learning the individual activity and aggregating them is computationally intensive and may not be informative compared to the collective context.

With the recent developments in deep learning research however, there has been a rise in developing models that can overcome some of these limitations. Notable among these techniques are the Convolution Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) for modelling spatial and temporal dependencies respectively. In this work, we focus on the task of recognising the collective movement activities of a group of sheep in a flock. We gather movement data of 36 sheep and investigate the use of several deep learning models and two features (spatial orientations and velocities) of the sheep to capture their collective activities. Our approach is one that learns in an online fashion the collective activity of the flock using a deep recurrent and convolutional neural network.

The main contributions of this paper are: (i) We extend the idea of activity recognition in animal behaviour to the collective settings and propose a model that can learn the collective movement behaviour in an online fashion (ii) We demonstrate that our approach can learn to classify activities even when some activities are underrepresented.

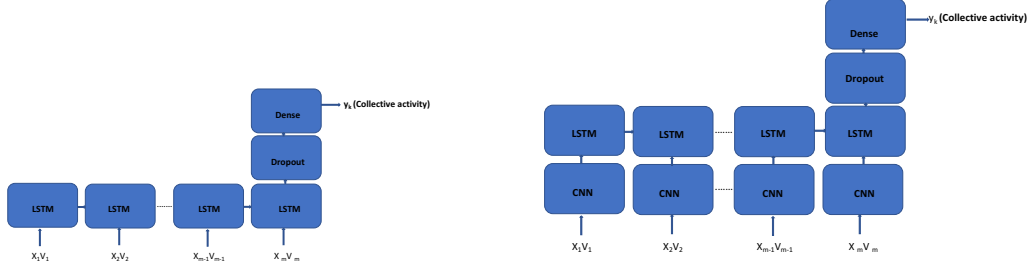


Figure 1: The basic LSTM and CNN+LSTM architectures on the left and right respectively. The inputs are $X_m V_m$ (orientation in space and velocity respectively) where m represents the size of the temporal window. Only the last output of both architectures is selected as the predicted activity (many to one architecture).

2 Problem Formulation

We briefly describe the formulation of the collective movement activity recognition problem here. Assume $D_{tr} = \{(X_n, V_n), y_k\}$ represents the training dataset of N_s samples and $D_{te} = \{(X_n, V_n)\}$ denotes the test set of N_d samples, where N_d may or may not be equal to N_s , and X_n, V_n represents spatial and velocity features of all the animals of interest respectively, while the corresponding label space representing the collective activities are $Y = \{y_1, y_2, \dots, y_k\}$ with k being the number of unique collective activities. Now the problem is as follows, given a history of observations over a time window t described by $\{X_i, V_i | i = 1, \dots, t\}$, the goal is to predict $y_k \in Y$ for each datapoint in D_{te} . That is, we aim to learn $p(y_{tk} | (X_t, V_t), (X_{t-1}, V_{t-1}), \dots, (X_{t-m}, V_{t-m}))$ where m is the time window of feature observations used for the prediction task.

3 Model Description

Temporal Dynamics Modelling: We use the recurrent neural network (RNN) to model temporal dependencies. In particular, we use the LSTM (Long Short-Term Memory) [6] a variant of the RNN to model long term dependencies and has been used in previous studies [7, 8] for handwriting synthesis and language translation. The LSTM can be described by the following equations:

$$\begin{aligned}
 i_t &= \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{ci}C_{t-1} + b_i) \\
 f_t &= \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{cf}C_{t-1} + b_f) \\
 C_t &= f_t C_{t-1} + i_t \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c) \\
 o_t &= \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{co}C_t + b_o) \\
 h_t &= o_t \tanh(C_t)
 \end{aligned} \tag{1}$$

where i_t , f_t , and o_t are input, forget, and output gate respectively. W is the weight matrix, x_t is the current input data, h_{t-1} is previous hidden output, C_t is the cell state and σ is the logistic sigmoid function.

Temporal and Spatial Dynamics Modelling: We use the recurrent convolutional network model of [9] where the input is passed through a CNN to produce a reduced representation of the input which is further passed into a RNN (Figure 1).

4 Data Acquisition and Pre-processing

Data collection: We present here our data collection system. All experiments involved with the animals complied with the Australian ethics laws guiding the handling of animals. We collected movement data of 36 sheep in a field with the aid of a GPS device designed in house for two days and attached to the back of the sheep. Previous work has shown that such harness equipment carried by sheep does not affect their locomotion [10]. Data was collected at a sampling rate of 1 sample/s to ensure all forms of interesting movement patterns by the sheep are captured. Each dataset contains a phase in which loggers were attached to each sheep in a holding pen followed by a phase in which the sheep were herded into the field and then a phase in which the sheep were left to roam across the field and a final phase in which the sheep were herded back to the holding pen to have the logging device removed and re-charged. We extracted only portions of the dataset where all the loggers were working. All missing data were interpolated between the last and next seen co-ordinates using the expectation-maximization algorithm [11].

Data labelling: With the aid of a viewer designed in house, we label the collective activities with respect to the instances where they occur in the dataset. Movement data of all the sheep in the flock were labelled with respect to the collective behaviours described in (Table 1). Due diligence was ensured to make sure the labelling was of high quality.

Table 1: Observed collective activities and their description.

Collective Movement Activities	Description
Not Active	Where the animals are gathered together in close proximity with little or no movement activities. Correspond to instances where the animals are resting or sleeping.
Active	The animals are moving and scattered in their habitat with diverse movement activities.
Herd Movement	The animals are being herded at a high velocity over a narrow space.

Table 2: Number of samples for each collective activity in the training and test datasets. The distribution can be seen to be highly skewed with the Herd movement mostly affected.

Dataset	Not Active	Active	Herd Movement
Train	21801 (37.55%)	35811 (61.68%)	452 (0.78%)
Test	26718 (41.96%)	36355 (56.24%)	597 (0.94%)

5 Experiments & Results

Experiments: We used dataset of activities for one day for training and another day for testing see (Table 2). Models were trained with the following parameters: Adam optimizer [12] at a learning rate = 0.001, look-back & recurrent cells = 30, batch-size=10 over 50 epochs, dropout = 0.2, categorical crossentropy for all losses and a softmax for all classification tasks, activation = tanh. For the CNN segment of the CNN+LSTM model, we used a 1D CNN with a 2x2 filter, stride=1 and relu activation. All models were basic with only one layer and trained on a 2.3 GHz Intel Core i5 PC. We investigate the use of two features, the velocities of the animals and the distance of each sheep to their centroid in an ablation manner with respect to the two models. The collective activities were one hot encoded. All evaluations were carried out with respect to the classification accuracy and confusion matrix.

Results: From the results in (Table 3), the LSTM (velocities & distance to centroid) architecture outperform others although marginally relative to the LSTM (velocities) and CNN+LSTM (velocities & distance to centroid) architectures. The distance of each sheep to the entire sheep centroid seems to be the least informative of the two features but produces better results when combined with the velocities of the sheep using LSTM. On the average the CNN+LSTM architectures perform slightly

better than the LSTM architectures over all feature combinations. More impressive is the fact that our model is able to classify very well the underrepresented activity. As seen in the confusion matrices (Figure 2), only the models including spatial features were able to learn the Herd movement activity with the CNN+LSTM performing better than the LSTM while the remaining two models¹ entirely misclassified this activity. This suggests that a fusion of spatial and movement features is essential to disentangle some of these complex collective activities especially in very rough and challenging terrains.

Table 3: Classification accuracy averaged over fifty epochs. Results show that the LSTM model with both spatial and velocity features gives better classification accuracy .

Models & Features	Classification Accuracy (%)
LSTM (velocities)	75.62 ± 0.25
LSTM (distance to centroid)	60.80 ± 1.91
LSTM (velocities & distance to centroid)	77.02 ± 1.11
CNN+LSTM (velocities)	73.65 ± 0.33
CNN+LSTM (distance to centroid)	70.63 ± 1.37
CNN+LSTM (velocities & distance to centroid)	74.51 ± 1.03

6 Conclusions & Future Work

In this paper, we have shown how to learn the collective movement activities of sheep using deep neural network. Our approach leveraged the fusion of spatio-temporal features from all animals of interest to learn collective activities even when their distribution is skewed. This work has implications for example in building automatic systems that can help farmers better understand the health of the flock as a whole for further applications in epidemic management as well as help conservationist learn collective behaviours of animals that are indicative of poaching activities. It is however not clear how our approach will perform when faced with a flock of different size. One potential solution is to use an embedding to project behaviours (features) into a low dimensional space similar to [13]. While we have used animals who live in co-operative societies, the method used here may not scale into other animal societies without any structure. In the future, we aim to explore the use of an hierarchical model to improve the classification accuracy most especially when the distribution of activities is skewed. While we have used very simple neural network architectures, in the future, it is our desire to investigate other complex deep learning architectures to improve the accuracy.

	Herd M.	Active	Not Active
Herd M.	0.84	0.16	0
Active	0	0.46	0.54
Not Active	0	0.03	0.97

(a) CNN+LSTM (velocities & dist. to centroid)

	Herd M.	Active	Not Active
Herd M.	0	1	0
Active	0	0.57	0.43
Not Active	0	0.13	0.87

(b) CNN+LSTM (velocities)

	Herd M.	Active	Not Active
Herd M.	0.73	0.27	0
Active	0	0.48	0.52
Not Active	0	0.03	0.97

(c) LSTM (velocities & dist. to centroid)

	Herd M.	Active	Not Active
Herd M.	0	1	0
Active	0	0.64	0.36
Not Active	0	0.15	0.85

(d) LSTM (velocities)

Figure 2: Confusion matrix for the top four performing models. The models where the spatial features were included show a significantly higher accuracy with respect to the underrepresented Herd movement activity.

¹We omit the remaining two models here for brevity.

References

- [1] Wongun Choi and Silvio Savarese. A unified framework for multi-target tracking and collective activity recognition. In *European Conference on Computer Vision*, pages 215–230. Springer, 2012.
- [2] Minsi Wang, Bingbing Ni, and Xiaokang Yang. Recurrent modeling of interaction context for collective activity recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [3] Pietro DeLellis, Giovanni Polverino, Gozde Ustuner, Nicole Abaid, Simone Macrì, Erik M Bollt, and Maurizio Porfiri. Collective behaviour across animal species. *Scientific reports*, 4: 3723, 2014.
- [4] Jacob W Kamminga, Helena C Bisby, Duc V Le, Nirvana Meratnia, and Paul JM Havinga. Generic online animal activity recognition on collar tags. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers*, pages 597–606. ACM, 2017.
- [5] Steffen Grünewälder, Femke Broekhuis, David Whyte Macdonald, Alan Martin Wilson, John Weldon McNutt, John Shawe-Taylor, and Stephen Hailes. Movement activity based classification of animal behaviour with an application to data from cheetah (*acinonyx jubatus*). *PLoS one*, 7(11):e49120, 2012.
- [6] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [7] Alex Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- [8] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [9] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
- [10] Hannah Hobbs-Chell, Andrew J King, Hannah Sharratt, Hamed Haddadi, Skye R Rudiger, Stephen Hailes, A Jennifer Morton, and Alan M Wilson. Data-loggers carried on a harness do not adversely affect sheep locomotion. *Research in veterinary science*, 93(1):549–552, 2012.
- [11] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pages 1–38, 1977.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Yongyi Tang, Peizhen Zhang, Jian-Fang Hu, and Wei-Shi Zheng. Latent embeddings for collective activity recognition. In *Advanced Video and Signal Based Surveillance (AVSS), 2017 14th IEEE International Conference on*, pages 1–6. IEEE, 2017.