Measuring Global Behaviour of Multi-Agent Systems from Pair-Wise Mutual Information

George Mathews[1], Hugh Durrant-Whyte[1], and Mikhail Prokopenko[2]

[1] ARC Centre of Excellence for Autonomous Systems
The University of Sydney, Sydney NSW 2006, Australia
[2] CSIRO ICT Centre, Locked Bag 17, North Ryde, NSW 1670, Australia
g.mathews@cas.edu.au

Abstract. In this paper we present a method for finding general dependencies between individual agents in a multi-agent system. We use mutual information as the general measure for dependence and demonstrate how it can be used to estimate pair-wise coupling in a 2 dimensional swarm. We then show that our technique of measuring local coupling can find phase transitions in the global behaviour of the system.

1 Introduction

When analysing complex multi-agent systems it is often hard to determine which components are coupled via the underlying system dynamics. We present a method of determining general nonlinear correlations between different objects that can be described by a vector of continuous variables. We use this local measure to analyse the global behaviour in a 2D swarm model.

Wright *et al.* showed in [1] that a dynamical systems approach can characterise the global behaviour of a swarm. The technique they developed required knowledge of the global state of the swarm. The technique we propose approaches the problem from an information theoretic perspective and looks for dependencies between the observable states of pairs of agents. The proposed metric captures the spatial and temporal coupling between the motion of the agents through state space. We then show that the mean and variance of these pair-wise dependencies can characterise different behaviours is swarming groups of agents.

The dependencies are measured by estimating their mutual information[2, 3]. Mutual information is introduced in Sect. 2, along with the method used to estimate it. In Sect. 3 we give an example of how mutual information can be used to calculate spatial and temporal coupling in a multi-agent system. In this section we also present a sampling technique to overcome issues such as non-stationarity and finite system run time, that is compatible with the mutual information estimator. We then give our conclusions and future work in Sect. 4.

2 Information Theoretic Measures of Dependance

Information theory provides many measures of dependence. Lets consider two continuous valued random vectors, $\mathbf{x} \in \mathbb{R}^{d_x}$ and $\mathbf{y} \in \mathbb{R}^{d_y}$, with marginal probability density functions, $p_{\mathbf{x}}(\mathrm{x})$ and $p_{\mathbf{y}}(\mathrm{y})$. The joint probability density function $p_{\mathbf{xy}}(\mathrm{x}, \mathrm{y})$, defined on the joint space $\mathbb{R}^{d_x d_y}$, is equal to the product of the marginal probability densities iff the two variables are independent. Any differences between $p_{\mathbf{xy}}(\mathrm{x}, \mathrm{y})$ and $p_{\mathbf{x}}(\mathrm{x})p_{\mathbf{y}}(\mathrm{y})$ are caused by some mutual relationship between the variables. It is here that information theory provides many different methods at comparing two probability density functions. The most common measure is the

Kullback-Leibler divergence[3] which corresponds to the Shannon mutual information when the joint probability density is compared to the product of the marginal densities. It is defined as

$$\begin{aligned} I(\mathbf{x};\mathbf{y}) &= D_{KL}(p_{\mathbf{xy}}(\mathrm{x},\mathrm{y})\|p_{\mathbf{x}}(\mathrm{x})p_{\mathbf{y}}(\mathrm{y})) \\ &= \iint p_{\mathbf{xy}}(\mathrm{x},\mathrm{y}) \log \frac{p_{\mathbf{xy}}(\mathrm{x},\mathrm{y})}{p_{\mathbf{x}}(\mathrm{x})p_{\mathbf{y}}(\mathrm{y})} d\mathrm{x}d\mathrm{y}. \end{aligned} \tag{1}$$

The base of the logarithm determines the units of the measure, for the work that follows we will assume base 2 and hence our mutual information is measured in bits.

2.1 Estimation of Mutual Information

One of the main difficulties with using mutual information (MI), especially in the continuous domain, has been trying to estimate it from sampled data. There are three main approaches used to estimate MI. The most popular are the histogram based methods[4, 5], these partition the continuous state state space into a number of *bins*. The probability densities are then estimated by the frequency counts of the bins. These methods often introduce large biases due to the discretisation. Another popular method relies on kernels to estimate the densities. This method is more theoretically sound for a finite number of samples, but relies on many tunable parameters and the final integration (in (1)) can be very difficult to compute. The final method is parametric based, which can only be used if a parametric model of the data already exists.

The method we will use here is most related to the histogram approach in the sense that the probability densities are approximated by sets of piece-wise constants, but it does not rely on specifying a partition on the state space. The MI is estimated by first estimating the marginal and joint differential entropies[2] of the variables

$$\hat{I}(\mathbf{x};\mathbf{y}) = \hat{H}(\mathbf{y}) + \hat{H}(\mathbf{y}) - \hat{H}(\mathbf{x},\mathbf{y}) \tag{2}$$

where $\hat{\cdot}$ refers to an estimate. Differential entropy differs from standard entropy since it is defined for continuous probability density functions, i.e.

$$H(\mathbf{z}) = -\int p_{\mathbf{z}}(\mathrm{z}) \log p_{\mathbf{z}}(\mathrm{z}) d\mathbf{z}. \tag{3}$$

This corresponds to the marginal entropy if $\mathbf{z} = \mathbf{x}$ or $\mathbf{y}$ and the joint entropy if $\mathbf{z} = [\mathbf{x},\mathbf{y}]^T$.

2.2 Differential Entropy Estimation

The method used here was developed by Kozachenko and Leonenko in [6]. This method has been reviewed in [7] and extended in [8]. For brevity, the full derivation of the estimator will be omitted, for a detailed proof the reader should see the above references.

It can be seen from (3) that differential entropy is proportional to the average of $\log p_{\mathbf{z}}(\mathrm{z})$, regardless of whether it is a marginal entropy or the joint entropy. Thus for a finite set of independent and identically distributed (i.i.d.) samples, $\{\mathrm{z}_i\}_{i=1}^n$, from $p_{\mathbf{z}}(\mathrm{z})$ we could estimate the entropy if we had a estimator for $\log p_{\mathbf{z}}(\mathrm{z}_i)$ by

$$\hat{H}(\mathbf{z}) = -\frac{1}{n}\sum_{i=1}^{n} \widehat{\log p_{\mathbf{z}}(z_i)}. \tag{4}$$

Based on this approach Kozachenko *et al.* proposed the estimator for the differential entropy based on the Euclidian distance, λ_i, between z_i and its nearest neighbour (i.e. $\lambda_i = \min \|\mathrm{z}_i - \mathrm{z}_j\|, \forall j \neq i$). For a given set of samples in $\mathbb{R}^d$ the estimator is given by

$$\hat{H}(\mathbf{z}) = \frac{d}{n}\sum_{i=1}^{n} E\{\log \lambda_i\} + \log\Big[\frac{S_d(n-1)}{d}\Big] + \frac{\gamma}{\ln(2)}. \tag{5}$$

Here, $E\{\log \lambda_i\}$ is the expectation of the logarithm of the distance to z_i's nearest neighbour, S_d is the surface area of a unit sphere in d-dimensional space, and γ is the Euler-Mascheroni constant ($= -\int_0^\infty e^{-\nu}\ln\nu d\nu \approx 0.5772156$). In general the surface area can be written as $S_d = \frac{d\pi^{d/2}}{\Gamma(r/2+1)}$ where $\Gamma(\cdot)$ is the gamma function.

The above estimator assumes that the density function is continuous and smooth, such that the density can be approximated by a constant in the vicinity of each sample point. To get a workable estimator a second approximation must be introduced, the expectation will be replaced with the actual observed value, i.e. $E\{\log \lambda_i\} \approx \log \lambda_i$. This estimator can now be used to obtain the marginal and the joint entropies required in (2) to estimate the MI.

3 Detecting Correlations in a Swarm

This section develops a method of using MI for detecting correlations in complex multi-agent systems. The system we will be analysing is a swarm model. We also show that the phase transition found by Wright *et al.* in [1] can also be found by our approach.

3.1 System Description

The system we will consider consists of N agents in $\mathbb{R}^2$. The state of agent i is specified by its position, $\mathrm{x}_i = [x_i, y_i]^T$, and velocity, $\dot{\mathrm{x}}_i = [\dot{x}_i, \dot{y}_i]^T$. Each agent has a sensor that can detect the relative position, r_{ij}, and velocity, $\dot{\mathrm{r}}_{ij}$, of all the other agents that are within its local neighbourhood (i.e. maximum sensor range). The agents also have an actuator that can apply an acceleration to the agent. The actuator is controlled with a onboard controller that takes as inputs the relative position and velocity of all the other observed agents and outputs a required acceleration. The controller implements two behaviours, an attraction/repusion law (ARL) based purely on the relative positions of the neighbouring agents and a damping law (DL) that is based on the relative velocity between the agent and its neighbours with a magnitude determined by the agents separation.

From these control laws the motion of each agent can be specified as

$$\ddot{\mathrm{x}}_i = \sum_{j\neq i} \boldsymbol{a}_{ARL}(\mathrm{r}_{ij}) + \sum_{j\neq i} \boldsymbol{a}_{DL}(\mathrm{r}_{ij}, \dot{\mathrm{r}}_{ij}). \tag{6}$$

The acceleration component from the ARL is directed toward the neighbouring agent j with a magnitude given in Fig. 1. The DL component is defined as $\boldsymbol{a}_{DL}(\mathrm{r}_{ij}, \dot{\mathrm{r}}_{ij}) = c_{DL}(\|\mathrm{r}_{ij}\|) \times \dot{\mathrm{r}}_{ij}$ where $c_{DL}(\|\mathrm{r}_{ij}\|)$ is a damping coefficient that varies as a function of the separation between agent i and j (see Fig. 1).

It can be seen from Fig. 1 that the controller is implicitly dependant on the neutral position of the ARL (i.e. r_{ARL}) and on the boundary of the DL (i.e. r_{DL}). If we fix the shape of these

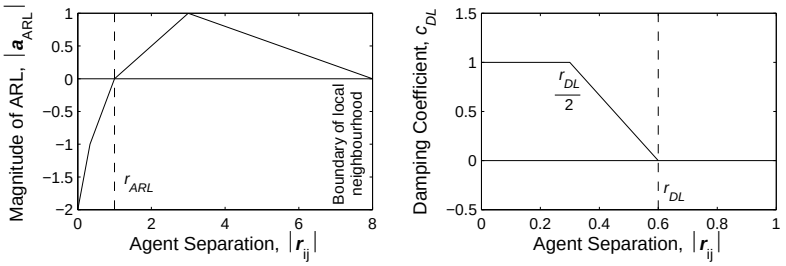

Fig. 1. LEFT: Magnitude of the ARL acceleration as a function of agent separation. RIGHT: Damping coefficient as a function of agent separation. For this current setting $\rho = r_{DL}/r_{ADL} \approx 0.6$

curves relative to r_{ARL} and r_{DL}, and scale all the distances by r_{ARL}, the global behaviour of the system can be specified by the single dimensionless parameter $\rho = r_{DL}/r_{ADL}$. Figure 2 shows typical system behaviours for different values of ρ. The system is initialised with a set of random positions and velocities, with all values evenly distributed on the interval $[-1, 1]$. The velocities are then modified such that the mean velocity, v_m, is always constant. This standardises the systems such that quantitative comparisons can be made that depend only on the parameter ρ.

3.2 Defining Correlations in a Multi-Agent System

Although a deterministic relationship exists between the full set of agents, we will examine how much coupling exists if we only consider a subset of the agents. The method we will pursue here relies only on drawing samples from the system and can be used for any system that has components with measurable outputs. But, this sampling limits us to only a few dimensions due to the computational complexity of the estimator.

Thus to estimate the dependencies in the system, we will calculate the MI between combinations of two agents and only consider their positions (i.e. ignore their velocities). That is, we want to know how much information is obtained about the position of one agent, $\mathbf{x}_i(t)$, at some unknown time t by knowing the position of another, $\mathbf{x}_j(t)$, at the same unknown time t. i.e. we want to compute

$$I(\mathbf{x}_i(t); \mathbf{x}_j(t)) \text{ for unknown } t \in [0, \infty). \tag{7}$$

This captures the spatial and temporal correlations between pairs of agents. By computing this for all combinations of agents we can build up a matrix of pair-wise correlations between the agents. This matrix would be similar to the covariance matrix but can be defined for objects that can consist of many variables and represents the total nonlinear dependance and not just linear dependance.

The difficulty in evaluating or estimating (7) comes from the fact that we cannot easily produce i.i.d. samples from the joint and marginal densities from a single instance of the system. This is due to the fact any measurement of the state of the system is highly dependent on the previous measurement. To overcome this, a monte carlo approach could be used to produce samples from many simulations of the system, but this will not be pursued since we cannot preserve the individual agents' identity between simulations and will prevent us from building up a pair wise correlation matrix for the group. Thus a different approach will be pursued that will involve transforming the system into a stationary one where we can draw i.i.d. samples from it and use them to estimate the MI of the original system.

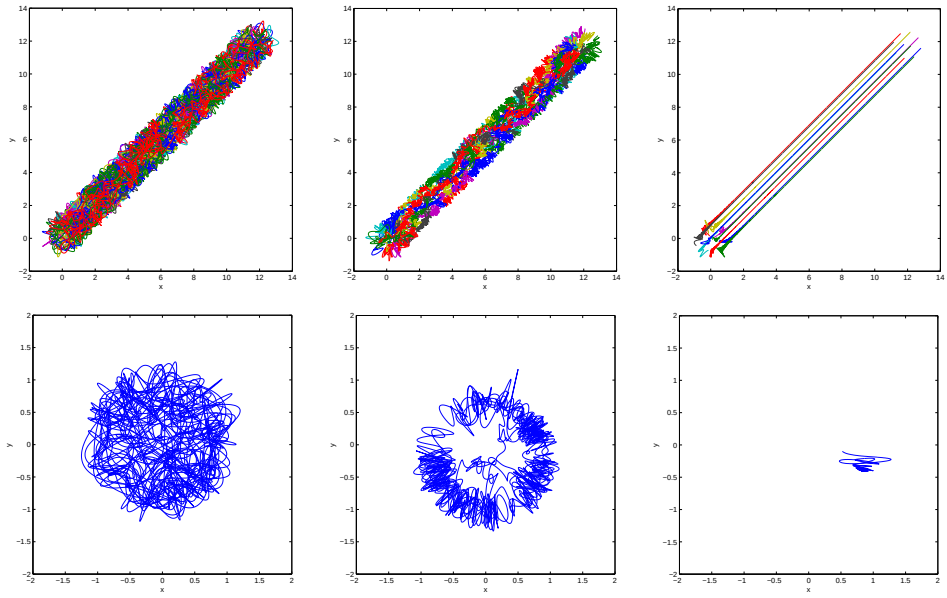

Fig. 2. Example trajectories for different values of ρ. For left frames $\rho = 0.28$, middle frames $\rho = 0.48$, and for right frames $\rho = 0.68$. The top row shows the trajectories of all 10 agents while the bottom shows the trajectory of a single agent relative to the centre of the group.

By creating a new coordinate frame centred on the mean position of the agents we can create a quasi-stationary system, in the sense that samples drawn from one instance of the system (i.e. at different times along the transformed trajectory) are representative of all instances of the system for any time. Trajectories from this are shown in the bottom row of Fig. 2. The transformation that produces these trajectories for agent k is

$$\mathbf{y}_k(t) = \mathbf{x}_k(t) - t\mathrm{v}_m \tag{8}$$

where v_m is the mean velocity of the system (which is fixed for all instances of all systems). In this coordinate frame, the probability density for the location $\mathbf{y}_k$ of agent k is $p_{\mathbf{y}}(\mathrm{y}_k)$ and is independent of time. From this density we can obtain the probability density function of the agent in the original coordinate frame at a specified time t, as $p_{\mathbf{x}}(\mathrm{x}_k|t) = p_{\mathbf{y}}(\mathrm{y}_k = \mathrm{x}_k - t\mathrm{v}_m)$. The density of $\mathbf{x}_k(t)$ for an unspecified time $t \in [0, \infty]$ can be recovered by noting $p_{\mathbf{x}}(\mathrm{x}_k) = \int p_{\mathbf{x}}(\mathrm{x}_k|t)p(t)dt$, where $p(t)$ is a uniform distribution on the interval $[0, \infty)$. This gives

$$p_{\mathbf{x}}(\mathrm{x}_k) = \lim_{T\to\infty} \frac{1}{T} \int_0^T p_{\mathbf{y}}(\mathrm{x}_k - t\mathrm{v}_m)dt. \tag{9}$$

This represents the probability density function for the location, $\mathbf{x}_k$, of agent k for any time $t \in [0, \infty)$ based on the probability density function of the agent's position relative to the centre of the group. We could try to obtain an expression for the MI based on this marginal probability density function (and a corresponding joint probability density function) but this would be impractical due to the limit extending to infinity. We can achieve the same aim (i.e. determining residual dependencies) by dropping the limit and taking it over a finite time horizon. From this theoretical analysis we will develop a method to produce i.i.d. samples from $p_{\mathbf{x}}(\mathrm{x}_k)$ by using the trajectories $\mathbf{y}_k(t)$.

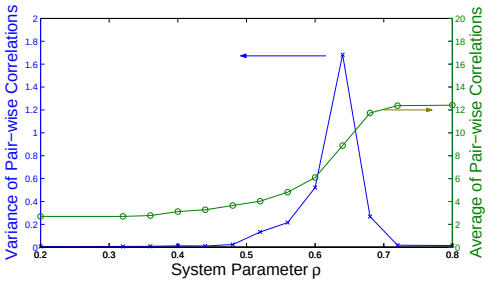

Fig. 3. Average and variance of the pair-wise mutual information for different systems.

Due to the quasi-stationarity of the trajectories, samples drawn randomly in time, from a trajectory, $\mathbf{y}_k(t)$, can be assumed to be i.i.d. from the density $p_{\mathbf{y}}(\mathrm{y}_k)$. These samples also correspond to i.i.d. samples from $p_{\mathbf{x}}(\mathrm{x}_k|t=0)$. To obtain samples from $p_{\mathbf{x}}(\mathrm{x}_k)$, where t is restricted to $[0,T]$, each sample is projected forward in time by a different random amount, $t^* \in [0,T]$. This projection simply corresponds to translating each sample by $t^*\mathrm{v}_m$.

Hence a sample, $\mathbf{x}_k^*$, from $p_{\mathbf{x}}(\mathrm{x}_k)$ can be obtained by a sample, $\mathbf{y}_k^*$, from $p_{\mathbf{y}}(\mathrm{y}_k)$ and a time t^* which is a random element of $[0,T]$, via $\mathbf{x}_k^* = \mathbf{y}_k^* + t^*\mathrm{v}_m$. The process is identical for obtaining samples, $(\mathbf{x}_i^*, \mathbf{x}_j^*)$, for the joint density, $p_{\mathbf{xx}}(\mathbf{x}_i, \mathbf{x}_j)$, from samples of $p_{\mathbf{yy}}(\mathbf{y}_i, \mathbf{y}_j)$. Thus using these samples, and the method outlined in Sect. 2, we can estimate the mutual information, $I(\mathbf{x}_i; \mathbf{x}_j)$, between all combinations of agents.

3.3 Results

Results were obtained for values of ρ from 0.2 to 0.8 for systems with 10 agents, with each estimate being produced with 8000 sample points. The results will be presented in two ways, the first will be the mean and variance of the pair-wise correlations. The mean characterises how structured the overall system is, while the variance provides information about sub-structure within the system. These results are shown in Fig. 3.

The second series of results that where obtained are the pair-wise information matrix, which contains the mutual information between all the combinations of agents. Some of these are shown in Fig. 4. The mutual information between an agent and itself has not been calculated and has been set to zero in the figure. It is noted that for small ρ and large ρ the size of the correlations is quite uniform across the system, where as in between (e.g. for $\rho = 0.64$) the system exhibits large differences between the correlations of the individual agents. This can also be seen in Fig. 3 where the variance of the pair-wise mutual information is also plotted.

This variations can be explained by looking at the dynamics of the system, for low ρ there is very little coupling between the agents and hence the correlation should be quite low. As the coupling increases (through increasing ρ and hence the damping), the system starts to undergo a phase transition, with local coupling between neighbouring agents. As ρ increases further the oscillations become completely damped out across the whole system and the system becomes globally coupled.

4 Conclusion and Future Work

In this paper we presented a method for finding general correlations between the motion of individual agents in a multi-agent system by estimating the mutual information between

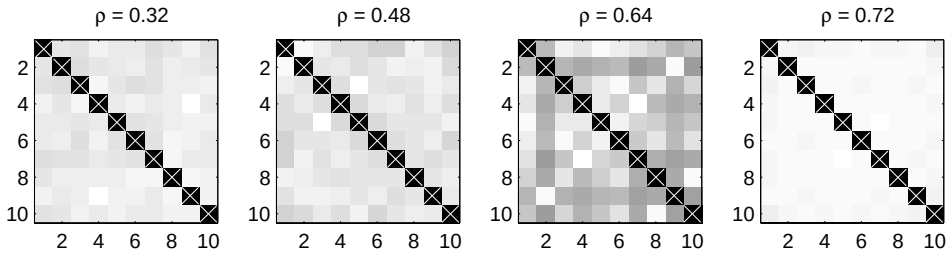

Fig. 4. Typical pair-wise mutual information matrices for different systems. White represents high correlation and black represents low correlation, with all values normalised by the maximum value for that system. The mutual information between an agent and itself has not been calculated.

them. We showed how this technique can be used to characterise the global behaviour of a simple multi-agent system. The important contribution of the work is to show that the global behaviour of the system can be characterised by examining only local pair-wise coupling. The global phase transition found by Wright *et al.* in [1] was also recognised by our technique.

For future work we hope to examine how this technique handles large systems with many components. For such systems it would be infeasible to calculate all pair-wise correlations, and we propose instead to only calculate the correlations for a subset of all the combinations across the system. We intend to show that this may be sufficient for a characterisation of global multi-agent behaviours.

Acknowledgements

This work is supported in part by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government, and in part by CSIRO Industrial Physics.

References

1. Wright, W.A., Smith, R.E., Danek, M.A., Greenway, P.: A measure of emergence in an adapting, multi-agent context. Proceedings of the Sixth Intenational Conference on the Simulation of Adaptive Behaviour (2000) 20–27
2. Cover, T., Thomas, J.: Elements of Information Theory. John Wiley (1991)
3. Kapur, J.: Measures of information and their Application. John Wiley (1994)
4. Moddemeeijer, R.: On estimation of entropy and mutual information of continuous distributions. Signal Processing **16** (1989) 233–248
5. Darbellary, G.: Estimation of the information by an adaptive partitioning of the observation space. IEEE Transactions on Information Theory **45** (1999)
6. Kozachenko, L., Leonenko, N.: A statistical estimate for the entropy of a random vector. Problems of Information Transmission **23** (1987) 9–16
7. Victor, J.: Binnless strategies for estimation of information from neural data. Physical Review E **66** (2002) 051903
8. Kraskov, A., Stogbauer, H., Grassberger, P.: Estimating mutual information. Physical Review E **69** (2004) 066138