Evaluating Team Performance at the Edge of Chaos

Mikhail Prokopenko, Peter Wang

CSIRO Mathematical and Information Sciences
Locked Bag 17, North Ryde, NSW 1670, Australia
{mikhail.prokopenko, peter.wang}@csiro.au

Abstract. We introduce a concise approach to teamwork evaluation on multiple levels — dealing with agent's *behaviour spread* and multi-agent *coordination potential*, and abstracting away the team decision process. The presented quantitative information-theoretic methods measure *behavioural* and *epistemic* entropy, and detect phase transitions — the *edge of chaos* — in team performance. The techniques clearly identify under-performing states, where a change in tactics may be warranted. This approach is a step towards a unified quantitative framework on behavioural and belief dynamics in complex multi-agent systems.

1 Introduction

The emergence of system-level behaviour out of agent-level interactions is a distinguishing feature of complex multi-agent systems — making them very different from other complicated multi-component systems, where multiple links among the components may achieve efficient interaction and control with fairly predictable and often preoptimised properties. In robotic soccer, the emergent behaviour is dependent on agents architecture and skills, the employed communication policy, the opponent tactics and strategies, and not least on various unknown factors present in the environment. In short, it appears to be extremely difficult to rigorously investigate and evaluate multi-agent teamwork, coordination, and overall performance. One possible avenue for measuring team performance is to use information-theoretic methods. In particular, we suggest to characterise dynamics of multi-agent teams in terms of generic information-theoretic properties, such as entropy, and correlate it with the overall team performance metrics.

Information-theoretic methods are applied in many areas exhibiting multi-agent interactions. For instance, Cellular Automata (CA) are a well-studied class of discrete dynamical systems, where information-theoretic measures of complexity (such as Shannon entropy of certain frequency distributions) were effectively used to categorise and classify distinct emergent configurations and phase transitions between them [17, 6]. Langton has shown in his seminal work [6] that an increase in the mutual information (defined as a function of individual cell entropies for a particular value of the λ parameter) is an indication of a phase transition from "order" to "chaos". Wuensche [17] has used a similar quantitative metric — variance of input-entropy over time — in classifying rule-space of 1-dimensional CA into ordered, complex and chaotic cases, related to Wolframs's qualitative classes of CA behaviour [16].

It could be argued that the complexity of emergent behaviour increases with a) the complexity of the agents, b) the diversity of the agents, achieved either by original design or by a learning process, and c) the variety of the communication connections among agents. In the context of RoboCup, the behavioural diversity (the second component of our argument) was extensively analysed by Balch [1], who suggested a new metric — hierarchic social entropy — to characterise the heterogeneity of agents behaviours across a team. Robots are said to be *absolutely behaviorally equivalent* if and only if they select the same behaviour in every perceptual state. Balch introduced the concept of hierarchical clustering as "a means of dividing a society into subsets of behaviorally equivalent agents at a particular taxonomic level" [1], and developed a measure of *behavioral difference*, enabling agent categorisation and subsequent calculation of the hierarchic social entropy.

We will initially focus on the first component — the diversity of a single agent's behaviour in different situations. In other words, we analyse a relation between entropy of an individual agent's behaviour and the team performance. Our conjecture, supported by experimental results, is that each agent is able to express more versatile behaviour when faced with easier opposition. Conversely, when opposing stronger teams, each agent may not be able to realise its behaviour in full — leading to lower *behavioural entropy*. Intuitively, these two extremes resemble the "ordered" and "chaotic" states: when the opponent is too strong then the agent's behaviour is limited to "fixed point" or "limit cycle" attractors, while weak opponents do not put significant constraints allowing the agent to achieve "strange" attractors symptomatic of chaotic behaviour. If this conjecture is true, then "complex" behaviour lies at *the edge of chaos*, and the behavioural entropy would point to a phase transition in this region. Put simply, when playing opponents of similar strength the agents exhibit most interesting "complex" behaviour, but at the same time it becomes much harder to evaluate the performance.

In the second part of the paper we study the third component — the complexity of the inter-agent communications, related to potential of multi-agent coordination. The analysis is focused on entropy of joint beliefs — the *epistemic entropy* — and complements the results reported earlier [10]. The epistemic entropy approach uses the information entropy as a precise measure of the degree of randomness in the agents' joint beliefs. Intuitively, the system with near-zero epistemic entropy (almost no "misunderstanding" in joint beliefs) has a higher multi-agent coordination potential than the system with near-maximal entropy (joint beliefs are almost random). In addition, we identified and considered two coupled levels of dynamic activity (following the Kugler-Turvey model) — showing that self-organisation and the loss of epistemic entropy occur at the macro (agent coordination) level, while the system dynamics on the micro level (within the communication space) generates increasing disorder. The entropy within the communication space is also traced against team performance metrics, showing that phase transitions occur in coordination-communication dynamics as well.

In summary, the developed metrics allow us to evaluate team performance on multiple levels: from individual "behavioural spread" to multi-agent coordination potential.

2 Input-Entropy and the Edge of Chaos

2.1 Mutual Information and Phase Transitions

The information-theoretic analysis of phase transitions is typically based on the notions of *entropy* and *mutual information*. The *entropy* is a precise measure of the amount of freedom of choice in the object — an object with many possible states has high entropy.

Formally, the entropy of a probability distribution $P = \{p_1; p_2; \ldots; p_m\}$ is defined by

$$H(P) = \sum_{i=1}^{m} p_i * \log(1/p_i) \tag{1}$$

The ratio of the actual to the maximum entropy is called the *relative entropy* of the source [14]. Langton [6] investigated the mutual information of CA, defined as a function of the individual cell entropies, $H(A)$ and $H(B)$, and the entropy of the two cells considered as a joint process, $H(A, B)$, that is: $I(A; B) = H(A) + H(B) - H(A, B)$, and related it to phase transitions. The average mutual information $I(A; B)$ has a distinct peak at the transition point: "the jump . . . clearly indicates the onset of the chaotic regime, and the decaying tail indicates the approach to effectively random dynamics".

Peaks or discontinuities are symptomatic of phase transitions in complex multi-agent systems. For instance, Miramontes [7] analysed artificial ant societies, composed of interacting agents that can generate regular cycles in the activity of the colony, and pointed out that the information capacity of the colony is maximal at certain nest densities — in the neighbourhood of a chaos-order phase transition. In other words, the maximum in the average information capacity of the colony, given by the classical Shannon entropy, corresponds to "the density at which the nest reaches its highest diversity of activity states". When the nest density is increased beyond some critical density and the phase transition has occurred, "the number of ants becomes sufficiently large to facilitate and support the existence of long-range correlated behaviour that manifests itself as coherent collective oscillations in the number active ants" [7].

Another way to identify phase transitions is to use *a variance of input-entropy*. Wuensche [17] characterised rule-spaces of 1-dimensional cellular automata with the Shannon entropy of rules' frequency distribution. More precisely, given a rule-table (the rules that define a CA), the input-entropy at time step t is defined as

$$S^t = -\sum_{i=1}^{m} \frac{Q_i^t}{n} \, log \frac{Q_i^t}{n} \ ,$$

where m is the number of rules, n is the number of cells (system size), and Q_i^t is the look-up frequency of rule i at time t — the number of times this rule was used at t across the CA. The input-entropy settles to fairly *low* values for ordered dynamics, but fluctuates irregularly within a narrow *high* band for chaotic dynamics. For the complex CA, order and chaos may predominate at different times causing the entropy to vary. A measure of the variability of the input-entropy curve is its variance or standard deviation, calculated over time. Wuensche has convincingly demonstrated that only complex dynamics exhibits high variance of input-entropy, leading to automatic classification of the rule-space. Importantly, the peak of input-entropy variance points to a phase transition again, indicating the edge of chaos (complexity).

2.2 Measuring Agent's Behavioural Spread

Having identified the appropriate metrics and the forces shaping the space-time dynamics, we now proceed to the analysis of the heterogeneity of a single agent's behaviour in different situations. As mentioned earlier, we shall explore a relation between the entropy of an agent's behaviour and the team performance. Our intention is to characterise an agent's behaviour without a loss of generality, and thus we would prefer to abstract away a possibly convoluted decision-making mechanism. In other words, we intend to consider only the action rules (condition-action pairs) that the agent triggered at a given time step. In other words, we may use (and did use) the agents designed in accordance with the Deep Behaviour Projection (DBP) agent architecture [10] or another (multi-layered) architecture, but without taking into account the depth of the agent behaviour in terms of multiple decision-making layers. This allows us to apply the developed techniques to any rule-based approach capable of identifying action rules taken by an agent at a given time step. To perform our analysis we employ the input-entropy of a particular frequency distribution B_i^k, where k is a game index, and i is an action rule index: $1 \leq i \leq m$, where m is the number of rules. Analogously to the CA analysis conducted by Wuensche [17], we define the behavioural input-entropy as

$$E^k = -\sum_{i=1}^{m} \frac{B_i^k}{n} \, log \frac{B_i^k}{n} \ ,$$

where n is the system size (the total number of rule invocations), and B_i^k is the look-up frequency of rule i during the game k. The difference between S^t and E^k is that the former is calculated for each temporal state of the CA in point, while the latter is determined for each game in a multi-game experiment. Both metrics, however, characterise the distribution of rules — either across the CA lattice or during the game.

We intend to show that agents express more diverse behaviour when faced with easier opposition. Formally, the average behavioural input-entropy, calculated for K games against the opponent j: $E_j = \sum_{k=1}^{K} \frac{E^k}{K}$, should in general increase with the average score difference g_j, defined as the average difference between the agent's team score and the opponent team score. Importantly, a standard deviation σ_j of the behavioural entropy E^k calculated across all games against the opponent j, will be shown to be an indicator of a phase transition, reaching a maximum at some g_j close to zero.

3 Experiments: Behavioural Entropy and Phase Transitions

We have carried out our experiments in the RoboCup Simulation League [4], where the platform simulates essentially a pseudo real-time environment, providing each agent with fragmented, localised and imprecise (noisy and latent) information about the environment. Each experiment included 30 games between the test team and a particular opponent, producing a value for the behavioural entropy E^k, $1 \leq k \leq 30$, and a score difference g_j (a negative score difference represents losing the game). Figure 1 shows input-entropy trajectories for 3 experiments, ranging from a much stronger opponent (the average score difference $g = -6.33$), to an opponent of about the same strength as the test team ($g = -0.07$), to a much weaker opponent ($g = +10.17$). It is easy to observe that not only the the behavioural entropy E^k of a test agent (the left mid-fielder of the test team, in this case) decreases on average with the strength of the opponent, but also that E^k fluctuates in a much wider band in the medium case.

To support this claim and to verify our conjecture that there is a phase transition, however, we conducted more experiments — against 10 opponents, collecting the statistics for 6 agents (wing- and centre-forwards, wing- and centre-midfielders, and wing- and centre-defenders). Figure 2 shows the average behavioural entropy (after $K = 30$ games), plotted for these 6 agents and for each of the opponents.

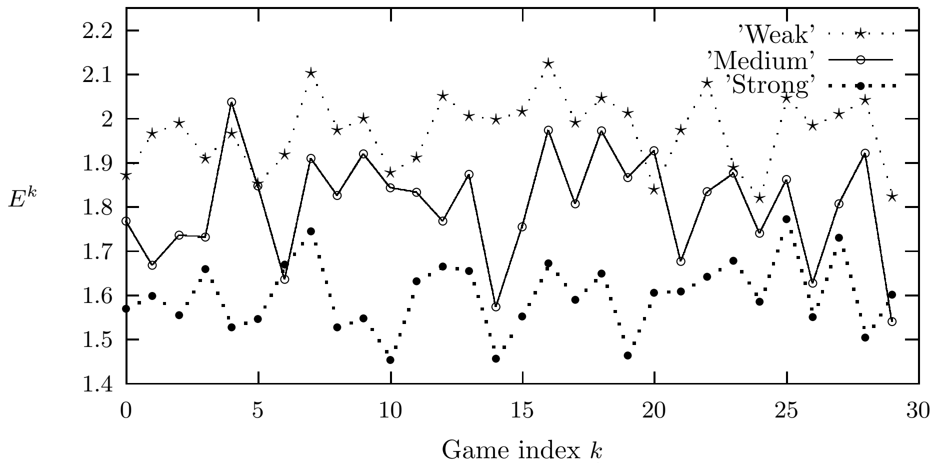

Fig. 1. Behavioural entropy E^k.

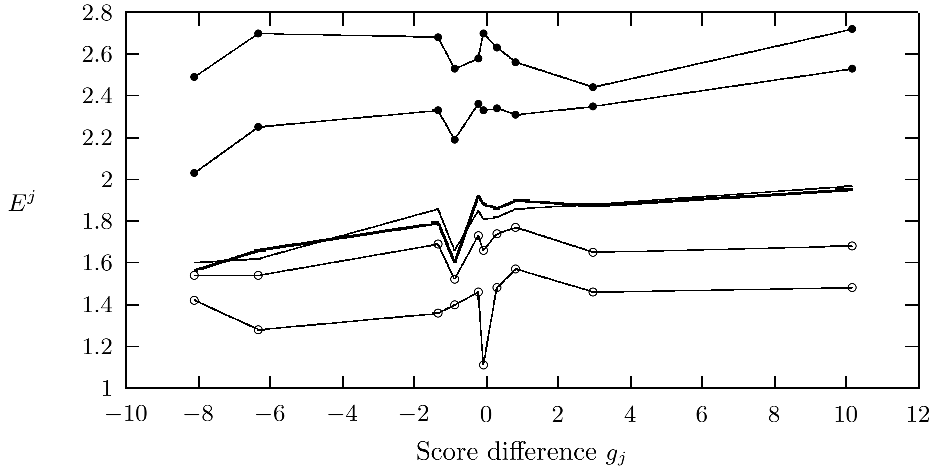

Fig. 2. Average behavioural entropy E_j. Two top plots represent forwards, two middle plots — midfielders, and two bottom plots — defenders.

The tendency of the behavioural entropy to increase when faced with a weaker opposition is obvious in most field positions, and especially in the midfield. There is also an evident discontinuity exactly in the range we expected — when competing with the opponents of similar strength ($-0.87 \leq g_j \leq 0.83$). This discontinuity is indicative of a phase transition. To confirm this, we observe the trajectory of standard deviation σ_j of the behavioural entropy E^k, calculated across all games against the opponent j, and shown in Figure 3. Standard deviation peaks in the expected region for all positions, and conclusively points to a phase transition. Interestingly, precise locations of peaks differ within the narrow range ($-0.87 \leq g_j \leq 0.83$), indicating that the peak is not a feature forced by a particular opponent, but rather a "complexity" attribute of the transition.

This entropy-based technique clearly identifies the *edge of chaos*. This is important because it helps to answer the question of whether a change in tactics is needed in some under-performing cases. During the phase transition, the team performance is unstable and the score difference should not be used as a sole trigger for drastic design changes or on-line interventions by a coach-agent.

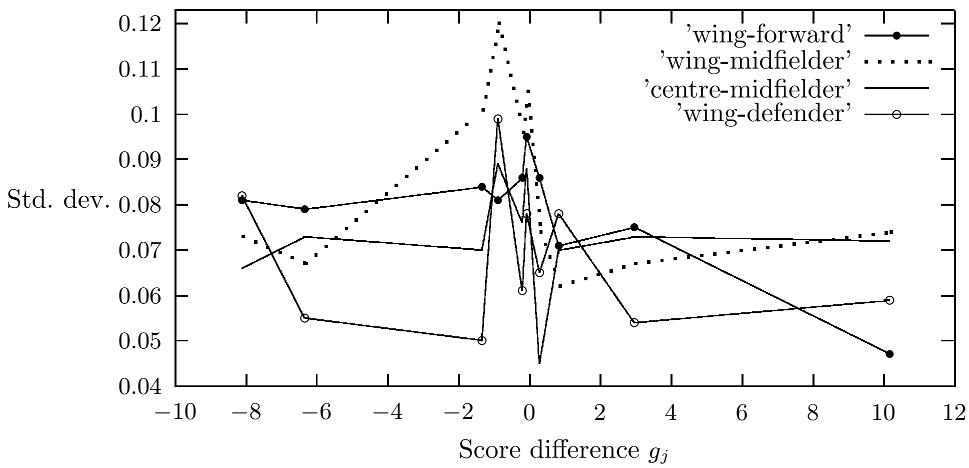

Fig. 3. Standard deviation of behavioural entropy.

4 Epistemic Entropy and Multi-Agent Coordination

4.1 Epistemic entropy on macro level

In this section we analyse the complexity of inter-agent communications. As pointed out in the literature [8, 5], emergent self-organisation or *extropy* may seem to contradict the second law of thermodynamics that captures the tendency of systems to disorder. The "paradox" has been gracefully explained in terms of multiple coupled levels of dynamic activity (the Kugler-Turvey model [5]) — self-organisation and the loss of entropy occurs at the macro level, while the system dynamics on the micro level generates increasing disorder. One convincing example is described by Parunak and Brueckner [8] in context of pheromone-based coordination. Their work defines a way to measure entropy at the macro level (agents' behaviours lead to orderly spatiotemporal patterns) and micro level (chaotic diffusion of pheromone molecules). In other words, the micro level serves as an entropy "sink" — it permits the overall system entropy to increase, while allowing self-organisation to emerge and manifest itself as coordinated multi-agent activity on the macro level. The epistemic entropy presented here is analysed in the terms of the Kugler-Turvey model [5] as well. We intend to show that the higher team coordination potential is related to lower entropy of multi-agent joint beliefs (macro level). At the same time, it is explained by increased entropy on a micro level. This micro level is the communication space where the inter-agent messages are exchanged (the process that is similar to diffusion of the pheromone molecules).

For convenience, we reproduce here definitions and two characterisations presented in [10], and follow with extended results. Let us consider a simple protocol $\mathcal{P}_1$ allowing an agent to communicate data about only one agent precisely. In other words, each agent is able to encode either the data about itself or about the other agent. Without loss of generality, we may assume that the protocol $\mathcal{P}_1$ has enough symbols to encode the data about agents (objects) $a_1, \ldots, a_n$ in such a way that they are explicitly distinguishable.

A binary relation $S(a_i, a_j)$ represents that the agent a_i *sends* a message containing the object a_j. A function C maps an agent name (symbol) to another agent name (symbol), and the abbreviation $C(a_i) = a_j$ denotes that the *content* of the message from the agent a_i is the object a_j. We intend that $C(a_i) = a_j$ if and only if $S(a_i, a_j)$.

Definition 1. *A multi-agent agreement $L_1(n)$ is called* selfish *if and only if $S(a_i, a_i)$ for all agents a_i, $1 \leq i \leq n$.*

A multi-agent agreement $L_2(n)$ is called transitively-selfish *if and only if $S^*(a_i, a_i)$ for all agents a_i, $1 \leq i \leq n$.*

Equivalently, $C(a_i) = a_i$ for all agents a_i, $1 \leq i \leq n$, under the selfish multi-agent agreement — each agent symbol is *a fixed-point* of the function $C(a)$. A transitively-selfish agreement suggests that the agents are more cooperative, and may communicate the data about some other agent (when available). Notice, however, that given the transitive closure $S^*(a_i, a_i)$, everyone is in the "loop". By definition, a selfish multi-agent agreement is transitively-selfish. The difference, however, may lie in the presence or absence of fixed-points $C(a_i) = a_i$. The transitively-selfish agreements without fixed-points, where each agent is *cooperative*: $C(a_i) \neq a_i$, will be of special interest. Non transitively-selfish agreements are called *mixed*.

Definition 2. *A multi-agent agreement $L_3(n)$ among $n \geq 2$ agents, where some agents are selfish and the others are cooperative, is called* mixed. *There are (αn) agents such that $S(a_i, a_i)$, and $(1 - \alpha)n$ agents such that $S(a_i, a_j)$ where $i \neq j$.*

The α parameter is called the team composition parameter.

The mixed agreement $L_3^1(n)$ where all cooperative agents provide information about the selfish team-mates is called the mixed agreement of the 1st kind.

The mixed agreement $L_3^2(n)$ where all cooperative agents are transitively communicating among themselves is called the mixed agreement of the 2nd kind.

In order to capture the distinction among selfish, transitively-selfish and mixed agreements in a formal information-theoretic setting we shall analyse the joint "output" of inter-agent communication at the end of each period of team synchronisation [10]. More precisely, we analyse joint beliefs represented by the sequence K_t of individual beliefs at time t: $K(a_i, a_j)$, where $1 \leq i \leq n$ and $1 \leq j \leq n$; the belief-function K is defined for each agent pair. In order to estimate how much information is contained in the whole team after a period of team synchronisation — as the team information progresses from K_t to $K_{t'}$ — we need to answer how much choice would be there if one were to describe $K_{t'}$. To do so we calculate the relative entropy H_r of $K_{t'}$. The following representation results for protocol $\mathcal{P}_1$ were reported in [10].

Theorem 1. *Selfish agreements attain minimal entropy. Transitively-selfish agreements without fixed-points attain maximal entropy asymptotically when $n \to \infty$.*

The first part of the theorem basically states that whenever agents agree to communicate the data about themselves only, they eventually leave nothing to choice, always maximising their joint beliefs. The intuition behind the second part is that the pair-wise "ignorance" of agents grows faster than the transitively-selfish agreement can cope with.

The next results for protocol $\mathcal{P}_1$ are the extensions produced for mixed agreements.

Theorem 2. *Mixed agreements of the 1st kind attain bounded epistemic entropy.*

Mixed agreements of the 2nd kind attain bounded epistemic entropy, and attain the epistemic entropy of mixed agreements of the 1st kind asymptotically when $n \to \infty$.

In other words, the lower limit is not 0, meaning that absolute order is never achievable regardless of the team composition or the number of agents, while the upper limit is not 1, so that absolute randomness is avoidable as well. Following [10] and interpreting the extended results, we would like to point out that the relative epistemic entropy of joint beliefs in multi-agent teams serves as a generic indicator of the team coordination potential. In general, the following series is established for the epistemic entropy:

$$H_r(L_1(n)) \ \leq \ H_r(L_3^2(n)) \ \leq \ H_r(L_3^1(n)) \ \leq \ H_r(L_2(n)) \tag{2}$$

while the respective coordination potentials follow the reverse dependency.

4.2 Epistemic entropy on micro level

The epistemic entropy may now be analysed in terms of the Kugler-Turvey model [5]. The higher coordination potential of the team following the selfish agreement with near-zero epistemic entropy can be explained by an increased entropy on the micro level — the communication space where the inter-agent messages are exchanged. Clearly, in the case of the selfish agreement the communication space is quite saturated, and the entropy on the micro level increases dramatically. On the contrary, the transitively-selfish agreement may use the communication channel(s) rather sparingly, resulting in a lesser increase of entropy on the micro level — while attaining near-maximal epistemic entropy on the macro level (joint multi-agent beliefs are almost random).

A characterisation of the micro level (the entropy "sink") can be obtained if one estimates the "regularity" of the communication space. In order to carry out this analysis we consider low-bandwidth domains requiring real-time response — in other words, environments where heavy communication among team agents is impossible. For example, we may consider Periodic Team Synchronization (PTS) domains introduced by Stone and Veloso [15] for pseudo real-time environments. However, our analysis is applicable to more generic domains as well — what is important is that the communication channel is limited. More precisely, a multi-agent domain should contain a parameter h determining how many messages can be received by each agent in a cycle. In particular, we are interested in capturing situations (*communication clashes*) where communication messages exceed "hear capacity" h in a given cycle, and measuring the average severity, spread and regularity of clashes. Let us introduce the following notation:

$\theta(a)$ is a function returning 1 if a boolean expression a is true, and 0 otherwise.
$\kappa_{i,j}$ is a function returning the number of communication messages received from the team i at cycle j;
$\delta_{i,j}$ is a boolean function returning $true$ if $\kappa_{i,j} > h$, and $false$ otherwise;

The average severity of clashes in the team i is given then by

$$M_i(h) \ = \ \frac{\sum_{j=1}^{m} \theta(\delta_{i,j}) \ \kappa_{i,j}}{m}$$

where m is the number of cycles, while regularity of the series $\kappa_{i,j}$ can be measured with the auto-correlation function of an integer delay τ:

$$\gamma_i(\tau) \ = \ \frac{\sum_{j=\tau+1}^{m} (\kappa_{i,j} - \overline{\kappa_i}) \ (\kappa_{i,j-\tau} - \overline{\kappa_i})}{\sum_{j=\tau+1}^{m} (\kappa_{i,j} - \overline{\kappa_i})^2} \ ,$$

where $\overline{\kappa_i}$ is the series average. The auto-correlation function is equivalent to the power spectrum in terms of identifying regular patterns — a near-zero auto-correlation across a range of delays would indicate high irregularity, while auto-correlation with values close to one indicate very high regularity. Some of this regularity is, however, spurious and is related to the severity of clashes. Therefore, we believe that a better approximation of the entropy on the micro level (communication space) may be given by the ratio

$$\xi_i(\tau, h) = \frac{M_i(h)}{\gamma_i(\tau)}$$

This new statistics attempts to capture how much *regularity* in the series is there *per communication clash*, and invert the measure. Our conjecture is that there is a dependency complementary to the dependency 2 over the range of possible values of τ, given some hear capacity:

$$\xi_{L_2}(\tau) \;\leq\; \xi_{L_3^1}(\tau) \;\leq\; \xi_{L_3^2}(\tau) \;\leq\; \xi_{L_1}(\tau) \tag{3}$$

The higher entropy on the micro level (*communication*) corresponds to the lower epistemic entropy on the macro level (*coordination*), and in turn to the higher coordination potential.

5 Experimental Results: Bounded Epistemic Entropy

Importantly, clear boundaries limiting the team coordination potential are related to particular communication policies. It is, however, not trivial to demonstrate these limits experimentally. First of all, the coordination potential can not be measured directly — it can only be realised in concrete multi-agent interactions. Moreover, the actual multi-agent coordination can be comprehensively evaluated only through the overall team performance over some period of time. Secondly, the coordinated activities corresponding to different communication policies would have to sufficiently differ in order to generate a *pronounced* difference in team performance.

In the RoboCup Simulation League, the "hear capacity" h determines how many messages can be heard by each agent in a cycle — for example, one message per cycle ($h = 1$). To measure coordination potential via team performance, we varied communication policies, while leaving all other factors (agents skills and tactics) unchanged. This focussed the experiment on the dependency between communication policies (and therefore, resultant joint beliefs) and the team coordination potential. In addition, we attempted to engineer, by varying the communication policies, very distinct types of coordinated activities, ranging from very local multi-agent coordination to rather global (zonal) one. For each type we also calculated the statistics $M_i(1)$ and $\xi_i(3, 1)$ — in order to estimate the corresponding entropy on the micro level.

We investigated three communication policies based on the protocol $\mathcal{P}_1$. The first policy ("Press") modelled the transitively-selfish agreement, with high relative entropy and very local coordination, enabling a pressing aggressive game. The second policy ("Zonal") followed the selfish agreement, with low relative entropy and very global coordination, enabling a passing non-aggressive game. The third policy ("Mix") was aimed at some mixture of local and global coordination, balancing predominantly pressing game with some passing chances — truly a mixed agreement with (anticipated) bounded relative entropy. The results are presented in the Table 1. The "Press" policy

Team	Goals	Points	$\gamma(3)$	$M(1)$	$\xi(3,1)$	Team	Goals	Points	$\gamma(3)$	$M(1)$	$\xi(3,1)$
vs “A”						vs “B”					
Press	31–91	65	0.553	0.067	0.12	Press	105–70	152	0.533	0.070	0.13
Zonal	18–107	50	0.657	0.224	0.34	Zonal	114–53	180	0.641	0.217	0.34
Mix	30–127	63	0.522	0.068	0.13	Mix	118–65	172	0.495	0.071	0.14

Table 1. Results against 2 benchmarks after 100 games for each test.

showed the best performance against the stronger benchmark “A”, while the “Zonal” policy was the worst. The results against the weaker opponent “B”, on the contrary, indicate that a team coordinated zonally (across wider spaces) perform better than the aggressive pressing team. In other words, these two communication policies lead to sufficiently different coordinated activities (local vs global) and generate a pronounced difference in team performance. This is important because, as mentioned earlier, we attempt to trace the effect of coordination potential — a capacity that is measurable only via a difference in produced results. The “Mix” policy achieved intermediate results. Importantly, this mixed policy was *within the boundaries* marked by the first two variants (and closer to the first one), as suggested by the relative epistemic entropy. As expected, the entropy $\xi(3,1)$ on the micro-level supported our hypothesis:

$$\xi_{Press}(3) \;\leq\; \xi_{Mix}(3) \;\leq\; \xi_{Zonal}(3)\;,$$

contrasting with the epistemic entropy inequalities, where the “Zonal” policy is close to the theoretic minimum and the “Press” policy reaches the maximum.

6 Experimental Results: Epistemic Entropy and Phase Transitions

We have also investigated phase transitions in the communication space. This investigation is at preliminary stages, but some promising results were obtained. Again, an experiment included 30 games between the test team (using the “Mix” policy) and an opponent from the pool of the same 10 opponents as in the behavioural entropy experiments. Each experiment produced a value for the communication entropy $\xi_k(3)$, $0 \leq k \leq 30$, and a score difference g_j. The standard deviation σ_j^{ξ} of the entropy $\xi_k(3)$, calculated across all games against the opponent j, is plotted in Figure 4. As expected,

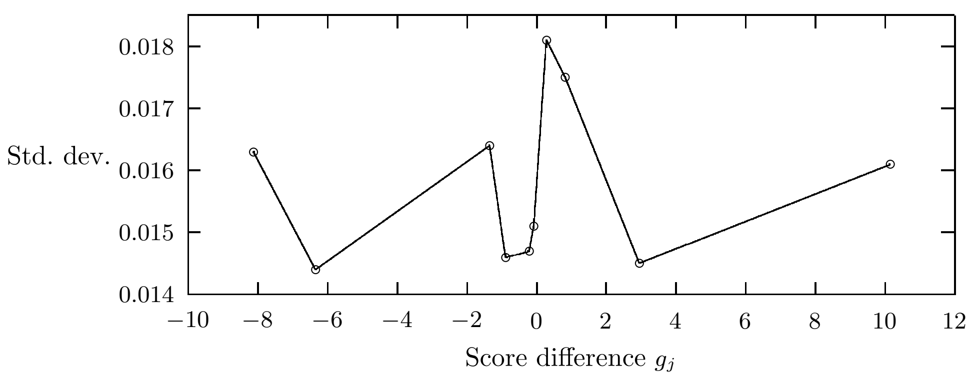

Fig. 4. Standard deviation of communication entropy.

standard deviation σ_j^{ξ} peaks in the the narrow range ($-0.87 \leq g_j \leq 0.83$), indicating a phase transition in the communication space, and in the coordination potential as well.

In other words, epistemic entropy, directly related to the entropy $\xi_k(3)$, also identifies the *edge of chaos*. In summary, the results not only illustrate the dependency between communication policy, the epistemic entropy and the team coordination potential, but also detect a phase transition in the coordination and communication dynamics.

7 Related Work and Conclusions

We presented a set of quantitative techniques for evaluation of team performance on multiple levels: from individual behavioural spread to multi-agent coordination potential. These techniques are based on information-theoretic metrics measuring complexity in multi-agent systems. In particular, we focussed on identifying the "edge of chaos" in team performance — leading to discovery of evident phase transitions. Our conjectures and theoretical results were supported by a number of experiments — over 500 games.

As pointed out by Pynadath and Tambe [11], "despite the significant progress in multiagent teamwork, existing research does not address the optimality of its prescriptions nor the complexity of the teamwork problem". The unified framework suggested by Pynadath and Tambe (COMmunicative Multiagent Team Decision Problem — COM-MTDP model) is general enough to subsume many existing models of multi-agent systems, and provides a breakdown of the computational complexity of constructing optimal teams in terms of observability and communication cost. The COM-MTDP model incorporates the team decision mechanism, and inevitably is rather complex, as almost any unifying framework. In this paper we attempted to introduce a concise approach to teamwork evaluation, dealing with *behaviour spread* and multi-agent *coordination potential*, and excluding the team decision process.

The presented analysis targets our overall goal — development of tools for evaluation of multi-agent adaptability and coordination in comparative terms, rather than methods for designing an "ultimate" intelligent and/or adaptive system. In pursuing this goal, we build up on existing quantitative methods for automated analysis of Simulation League games (eg., the AGL tool — Analysis of Game Logs [2]). We also hope to complement existing teamwork models and techniques impacting the team performance. A pioneering system capable of an automated analysis in the context of RoboCup was the ISAAC system modelling an assistant-coach [12]. ISAAC analyses games off-line and produces structured suggestions to designers, supported by examples. Another autonomous coach agent is recently described by Riley et.al. [13] — it is not only capable of extracting models from past games but may also respond to an ongoing game.

Our quantitative analysis complements this line of research by providing methods and techniques that determine phase transitions in team performance — the *edge of chaos*. These techniques isolate the under-performing cases where a change in tactics is warranted. During the phase transition, the team performance is highly unstable, and the scope for an on-line coach-agent contribution is limited. The quantitative information-theoretic methods presented here incorporate both behavioural and epistemic entropy, and are compatible with the hierarchic social entropy approach developed by Balch [1]. This opens a clear way to a unified quantitative framework on *behavioural and belief dynamics* in multi-agent systems. Another interesting possibility is to explore possible connections between the described techniques and the measures for relevant information investigated by Polani et.al. [9], as well as the conditions reducing "the influence of cognition difference" introduced by Cai et.al. [3] in the context of RoboCup.

Acknowledgements. The authors are grateful to members of the "Entropy and self-organisation in multi-agent systems" discussion group, and in particular to Mark Foreman, Ying Guo, Andrew Lampert, and Philip Valencia for many open and motivating discussions; and to Marc Butler, Thomas Howard and Ryszard Kowalczyk for their exceptionally valuable prior contributions to RoboCup Simulation efforts.

References

1. Tucker Balch. Hierarchic Social Entropy: An Information Theoretic Measure of Robot Team Diversity. *Autonomous Robots*, vol. 8, no. 3, July, 2000.
2. Marc Butler, Mikhail Prokopenko and Thomas Howard. Flexible Synchronisation within RoboCup Environment: a Comparative Analysis. In P. Stone, T. Balch, G. Kraetzschmar (eds.). *RoboCup-2000: Robot Soccer World Cup IV*, LNCS 2019, Springer, 2001.
3. Yunpeng Cai, Jiang Chen, Jinyi Yao, Shi Li. Global Planning from Local Eyeshot: An Implementation of Observation-Based Plan Coordination in RoboCup Simulation Games. In A. Birk, S. Coradeschi, S. Tadokoro (eds.). *RoboCup 2001: Robot Soccer World Cup V*, LNCS 2377, Springer 2002.
4. Hiroaki Kitano, Milind Tambe, Peter Stone, Manuela M. Veloso, Silvia Coradeschi, Eiichi Osawa, Hitoshi Matsubara, Itsuki Noda and Minoru Asada. The RoboCup Synthetic Agent Challenge. In Proceedings of the 15th Int'l Joint Conference on Artificial Intelligence, 1997.
5. Peter N. Kugler and Michael T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement.* Lawrence Erlbaum, 1987.
6. Christopher Langton. Computation at the Edge of Chaos: Phase Transitions and Emergent Computation. In S. Forest (ed.), *Emergent Computation*, MIT, 1991.
7. Octavio Miramontes. Order-Disorder Transitions in the Behavior of Ant societies. *Complexity*, Vol. 1, no. 3, 56-60, 1995.
8. H. Van Dyke Parunak and Sven Brueckner. Entropy and self-organization in multi-agent systems. In Proc. of the 5th Int'l Conference on Autonomous agents, Montreal, 2001.
9. Daniel Polani, Jan T. Kim and Thomas Martinetz. An Information-Theoretic Approach for the Quantification of Relevance. In J. Kelemen and P. Sosik (eds.), *Advances in Artificial Life*, Proceedings of the 6th European Conference on Artificial Life, Prague, Springer 2001.
10. Mikhail Prokopenko and Peter Wang. Relating the Entropy of Joint Beliefs to Multi-Agent Coordination. In Proceedings of the 6th International Symposium on RoboCup, 2002.
11. David V. Pynadath, Milind Tambe. Multiagent teamwork: analyzing the optimality and complexity of key theories and models. In the Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2002, Bologna, 2002.
12. Taylor Raines, Milind Tambe, Stacy Marsella. Automated assistants to aid humans in understanding team behaviors. In the Proceedings of the Fourth International Conference on Autonomous Agents, Barcelona, 2000.
13. Patrick Riley, Manuela Veloso, Gal Kaminka. An Empirical Study of Coaching. In H. Asama, T. Arai, T. Fukuda, and T. Hasegawa (eds.), Distributed Autonomous Robotic Systems 5, Springer-Verlag, 2002.
14. Claude E. Shannon and Warren Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.
15. Peter Stone and Manuela Veloso. Task Decomposition, Dynamic Role Assignment, and Low-Bandwidth Communication for Real-Time Strategic Teamwork. In *Artificial Intelligence*, vol. 110, no. 2, June 1999.
16. Stephen Wolfram. Universality and Complexity in Cellular Automata. *Physica D*, 10 (January 1984), 1–35, 1984.
17. Andrew Wuensche. Classifying Cellular Automata Automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the Z parameter. *Complexity*, Vol. 4, no. 3, 47-66, 1999.