

Quantifying the impact of communication on performance in multi-agent teams

Mathew Zuparic¹ · Victor Jauregui² · Mikhail Prokopenko² · Yi Yue¹

Received: 14 December 2016 / Accepted: 6 April 2017 / Published online: 1 June 2017
© Her Majesty the Queen in Right of Australia 2017

Abstract In this work, we relate the extent and quality of inter-agent communication and the overall performance in teams of multiple agents. Specifically, we examine the RoboCup Soccer Simulation 2D League, and carry out multiple simulation experiments against two evenly matched teams. For each simulated run (a 2D soccer simulation game), we generate the communication efficiencies (i.e., communications sent/communications received) for each agent pair. Applying linear regression and principal component analyses, we then correlate these efficiencies with measures of performance (i.e., goals scored and goals conceded), enabling the construction of inter-agent communication networks. Analysis of these networks highlights the microscopic player-to-player and macroscopic role-to-role communications correlated with performance. The approach determines the salient pathways within inter-agent communications which globally affect the coordination and the overall performance in multi-agent teams.

Keywords RoboCup · Multi-agent · Communication · Regression · Network

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 pre-optimised properties [57]. In many multi-agent setups, including command-and-control scenarios, the emergent behaviour is dependent on agent architecture and skills, the employed communication policy, the opponent tactics, and strategies, and not least on various unknown factors present in often adversarial environment [33].

Typically, it appears to be extremely difficult to rigorously investigate and evaluate multi-agent teamwork, coordination, and overall performance. One possible avenue for predicting team performance is to measure communication efficiency within a team during a scenario and estimate its impact on the team's performance. In particular, we intend to characterise coordination of multi-agent teams in terms of their communication efficiencies, suitably defined for multiple situations, message types, and contexts, and correlate these with the overall team performance metrics.

1.1 RoboCup domain

In pursuing our objective, we examine the domain of the RoboCup Soccer Simulation 2D League, which allows us to run multiple simulation experiments while varying a number of variables associated with agent communications. RoboCup (the “World Cup” of robot soccer) was first proposed in 1997 as a standard problem for the evaluation of theories, algorithms, and architectures for Artificial Intelligence, robotics, computer vision, and several other related areas [39], with the overarching RoboCup goal of developing a team of humanoid robots capable of defeating the FIFA World Cup champion team (the “Millennium Challenge”). Since 1997, The

✉ Mathew Zuparic
Mathew.Zuparic@dsto.defence.gov.au

¹ Decision Sciences, Defence Science and Technology Group, Department of Defence, Canberra, ACT, 2600, Australia

² Complex Systems Research Group, Faculty of Engineering and IT, The University of Sydney, Sydney, NSW, 2006, Australia

RoboCup initiative has convincingly superseded chess as a benchmark for Artificial Intelligence, developing along two general complementary paths [38]: physical robot leagues, and software agent (simulation) leagues [49].

RoboCup 2D Soccer Simulation League specifically targets the research question of how the optimal collective dynamics can result from autonomous decision-making under noisy conditions and several constraints, set by tactical plans and teamwork (collaboration) as well as opponent (competition) [15, 40, 46, 53, 56, 57, 61, 62, 71, 73, 75, 82]. It encourages development of diverse player behaviours and team strategies [6, 8, 35, 41, 54, 60, 70, 85], and offers a robust framework for evaluating the emergent collective behaviours and team performance.

In particular, RoboCup Soccer Simulation League (both 2D and 3D) involves software agents playing games on a centralised server (maintaining the “world model”, including player and ball dynamics and kinematics) over a network [17], and offers several obvious advantages in comparison to physical leagues, including the ability to simulate soccer matches without physical robots and abstracting away low-level hardware and environmental issues (e.g., motor temperature and breakages) [14]. As pointed out by Budden and Prokopenko [14], the simulation leagues often serve as platforms for the initial development and evaluation of software modules for later integration into physical robots [44, 48], and have applications well beyond the RoboCup domain (e.g., localisation and mapping [12, 30]).

Each team consists of 11 players and a “coach”, which is a non-playing agent responsible for assigning each agent a specific type, given a number of randomly generated physical profiles (including characteristics such as speed and stamina). Each of the fully autonomous simulated *agents* interacts with the soccer server, receiving information from the server relative to its current field of view; determining what actions to execute; and submitting these requests to the server. The server fulfills these requests and resolves any resultant conflicts (e.g., two agents attempting to occupy the same spatial location). The server proceeds in real time and imposes noise on both the agents’ observations and actions [49]. It is the responsibility of each agent to submit its action requests at the appropriate times to stay synchronised with the soccer server. Furthermore, each agent is allocated an individual CPU process, with no direct inter-process communication permitted. In addition, the soccer server provides a low-bandwidth, indirect communication method between the agents by supporting simulated verbal commands.

Crucially, although simulation league agents have only noisy perception of their environment, the soccer server itself has perfect information regarding the global state, enabling replicable quantification of experimental

performance, based on the ability to run massively-parallelised experiments [13, 14].

1.2 Multi-agent communication and coordination: related work

Analysis of various links between cooperation, coordination, and performance has been a continuing focus of research in multi-agent systems, within the general class of problems in which multiple agents have to coordinate their strategies to cooperate on some task. The term *team* generally refers to a set of agents with a common goal.

Approaches to activity coordination in multi-agent teams range from strictly top down (plan-based coordination) to purely emergent (reactive coordination), with many hybrid variants, each having its specific advantages and disadvantages. This diversity is directly affected by specific multi-agent communication policies which facilitate coordination and/or cooperation among agents.

For example, Stone and Veloso [72] investigate cooperation in multi-agent systems to improve team performance in the context of multi-agent learning. Several surveys of the literature on multi-agent systems are available [34, 74, 81]. Shoham and Leyton-Brown [68] visit the foundations of multi-agent systems.

More recently, as the ability to generate and collect data improved, there has been growing interest into using data-analytical methods in applied multi-agent domains. For example, Rein and Memmert [59] advocate the relevance and applicability of data analytics to tactical analysis of team sports. In a different domain, Ajitha et al. [3] apply multi-agent systems principles to software systems, with an emphasis on measuring cooperation among software agents.

There have also been various studies which investigate how communication in multi-agent systems contributes to improved performance of teams of agents. Communication has the potential to improve team performance in domains in which individual agents are only able to observe a small part of the world; the assumption of *partial observability* described by Veloso et al. holds. In such domains, agents can communicate to share information about their surroundings with other agents, thereby helping all agents compile more complete knowledge about the environment. Of course, communication usually comes at some cost. Bernstein et al. [10] study the computational complexity of communication in decentralised partially observable Markov decision processes (DEC-POMDPs).

The potential for communication to improve teamwork in multi-agent systems is discussed in Stone and Veloso’s survey of multi-agent learning [74]. As highlighted therein, communication can improve team performance in multi-agent systems, but has the draw-back of

increasing coordination complexity: agents must decide *when* and *what* to communicate to coordinate their activities towards achieving a team goal. Roth et al. examine heuristics for deciding when to communicate [63] and what to communicate [64] in environments where there is a cost associated with communication.

The application domains (e.g., RoboCup) considered in our investigation are compatible with the *COM-MTDP* model developed by Pynadath and Tambe [58]. This model makes communicative acts (viz., saying something) explicit, and distinguishes them from other action types. This is particularly pertinent in our analysis, as the effect of explicit communication is the main object of study.

Panait and Luke [50], in their survey of the space of multi-agent learning, explore scenarios in which agents benefit by communicating to solve a collective problem, viz., cooperative multi-agent learning.

Becker et al. [9] also investigate the question of when to communicate in domains in which there is a cost associated with communication. They define the net performance gain of communication by measuring the *value of communication* (VoC) as the difference between the performance of a team when agents communicate explicitly and the performance without communication. This approach follows Howard's [32] method of quantifying the value of information in decision problems. According to Becker et al., agents should engage in communication only when the VoC is such that it is expected to improve net team performance.

Gutiérrez et al. [28, 29] study metrics measuring the performance of multi-agent systems. Performance is measured by quality of service (QoS) as system responsiveness. The approach is based on load-balancing in communication networks and concentrates on identifying patterns which unevenly balance the communication load in a network of agents. QoS, as measured above, is measured as the delay in message delivery. This is less of a concern in domains such as ours in which, by contrast, uncoordinated communication can lead to message collisions, which results in some messages being lost. In these domains, the proportion of sent messages which arrive safely is an important measure.

Chou et al. [19] and Nair et al. [47] study communication and coordination as applied to emergency response settings. The latter explore how dynamic reconfigurability of teams can contribute to improved performance in environments that require decentralised control. Our domain is more dynamic in the sense that we have to deal with not only partial information, but there are other agents in active opposition. Wu et al. [83] have studied the effect of communication on decentralised on-line planning in multi-agent systems which act in highly dynamic environments. The domains that they

study share environmental properties that are common to ours.

In domains similar to ours, Candea et al. [16] explore the effects of coordination on team performance in simulated football. Communication is central to their approach. They express their ultimate goal as:

exploiting the use of communication among the players to improve team performance, allowing the robots to acquire more information, and to self-organize in a more reliable way.

—Candea et al. [16] (p80) Other recent related works in the RoboCup domain include Bai et al. [7] who use the celebrated *WrightEagle* team to show the scalability of the MAXQ-OP algorithm for effective multi-agent planning and decision-making under uncertainty. In addition, Hausknecht et al. [31] detail a platform for algorithm experimentation in multi-agent learning based on the *Half Field Offence* subtask.

Multi-agent coordination potential can also be quantified indirectly [56, 57], by characterising various inter-agent communication policies in terms of generic information-theoretic properties. Specifically, the complexity of the inter-agent communications has been related to the potential of multi-agent coordination by estimating the epistemic entropy [57] 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). Finally, the entropy within the communication space has been traced against team performance metrics, showing that phase transitions occur in coordination-communication dynamics as well [57].

There are other frameworks relevant to our application domain which use radically different methodologies. For an investigation into communication in robotic soccer using an argumentation-based framework, see Frias-Martinez et al. [25]. A general information-theoretic treatment of optimal inter-agent communication and the communication efficiency is provided by Prokopenko et al. [55] and Salge et al. [65].

1.3 Multi-agent team networks

Quantitative analysis, including network science methods, is increasingly being used in team sports to better understand and evaluate performance [1, 78]. For instance, Fewell et al. [24] analysed basketball games as strategic networks, where players are represented as nodes and passes as edges: the resulting network captures ball movement, at different

stages of the game. Their work studies network properties (degree centrality, clustering, entropy and flow centrality) across teams and positions (roles), and attempts to determine whether differences in team offensive strategy can be assessed by their network properties.

The study of Peña et al. [52] constructed a static-weighted-directed graph for each team (*the passing network*), using passing data made available by FIFA during the 2010 World Cup, with vertices corresponding to players and edges to passes. This provided a direct visual inspection of a teams strategy and determined the relative importance of each player in the game, using different centrality measures..

Recently, Cliff et al. presented several information-theoretic methods of quantifying dynamic interactions in soccer games, using the RoboCup 2D simulation league as an experimental platform [22, 23]. These interactions were detected information theoretically and captured in two ways: via (i) directed networks (interaction diagrams) representing significant coupled dynamics of the players positional data, and (ii) state-space plots (coherence diagrams) showing coherent structures in Shannon information dynamics.

In a general sense, the problem of constructing a network, given some (partially) observed dynamics, is related to the structure learning problem for spatially distributed dynamical systems [37]. As pointed out by Boccaletti et al. [11], modelling a partially observable system as a dynamical network presents a significant challenge in synthesising these models and capturing their global properties. There are many practical problems in this class, related to inferring a specific network structure, e.g., effective networks in neuroscience [42, 43, 51, 66, 69, 77], multi-agent systems [26, 84], dynamical Bayesian networks [18, 21, 27], among others. The prominent feature in these approaches is a consideration of both the *structural* and *functional* connectivity, and the inference of the functional topology based on underlying dynamics partially observed from distinct structural nodes.

In Sect. 2, we describe the methods used in this study, including communication data generation through multi-agent RoboCup simulation, and detail statistical analysis techniques applied to the data to infer communication (*functional*) networks based on the performance of a (*structured*) team of multiple agents. In Sect. 3, we present the communication networks which result from performing statistical analysis on the simulation data and analyse these results. Finally, in Sect. 4, we comment on further work which can eventuate from the presented results.

2 Methodology

Consider a domain in which a set of agents A communicate with each other by sending messages each of which may be categorised according to some set of message types M . Let

S be a set of situation types, such that an agent $a \in A$ may send a message of type $m \in M$ in a situation of type $s \in S$.

Let the set of *communication contexts* (or just *contexts* for short), C , with respect to S and M be the set of all situation-message pairs: i.e., $C = S \times M$.

2.1 Communication efficiency

2.1.1 General framework

For each context c , and every possible tuple of agent pairs (a, b) , we associate a real-valued index $e \in \mathbb{R}$ which measures the *communication efficiency* from agent a to agent b , with respect to c . The task then is to define for each context c a real-valued partial function $\varphi_c : A \times A \rightarrow \mathbb{R}$ which assigns to pairs of agents a and b a real number $e = \varphi_c(a, b)$ representing the efficiency of communication from the first agent to the second.

2.1.2 Domain refinement

For the domains considered in this paper, the set of agents, A , and the set of message types, M , are finite. Suppose further that messages may be sent only in any of some finite number of discrete time-steps/cycles, and that an agent may send at most one message per cycle. Moreover, suppose that the communication channels are unreliable in the sense that messages between agents may get lost, but are otherwise transmitted without noise. Under this restriction, only finitely many messages may be sent during the period under consideration.

Let $s_c(a)$ be the total number of messages sent by agent a during a given period of time, and let $r_c(a, b)$ be the number of those messages sent by a which were received by agent b over the same period. Provided agent a has sent at least one message (i.e., $s_c(a) \neq 0$); the communication efficiency $\varphi_c(a, b)$ from agent a to agent b is defined by the following:

$$\varphi_c(a, b) = \frac{r_c(a, b)}{s_c(a)}. \quad (1)$$

In other words, the communication efficiency from agent a to b is the proportion of messages sent by a which were received by b over a given period of time.

For example, if in a certain context over a given period of time agent a sent 25 messages, and agent b only received 18 of those messages, then $r_c(a, b) = 18$ and $s_c(a) = 25$. The communication efficiency during that period is $\varphi_c(a, b) = \frac{18}{25} = 0.72$.

Observe that, in general, $r_c(a, b) \neq r_c(b, a)$, and hence, φ_c is not symmetric; i.e., in general, $\varphi_c(a, b) \neq \varphi_c(b, a)$. Moreover, because each message received by an agent must have been sent by another, it follows that

$0 \leq r_c(a, b) \leq s_c(a)$. Therefore, for any context \mathbf{c} and pair of agents a, b , it follows that $\varphi_c(a, b) \in [0, 1] \subseteq \mathbb{R}$.

2.2 Network representation

A system of communicating agents may be represented as a weighted network:

$$\mathcal{N} = (V, E, w)$$

where V is a set of nodes, E is a set of directed links between nodes ($E \subseteq V \times V$), and $w : E \rightarrow \mathbb{R}$ is a function associating a weight to each link. A link between two agents represents the communication between the two.

2.2.1 Communication networks

A *communication network* is defined to be a network in which the nodes are agents, i.e., $V = \mathbf{A}$. Moreover, a communication network includes no reflexive links; i.e., for each node, $v \in V$, $(v, v) \notin E$. This restriction reflects the intuition that agents do not send messages to themselves. For each context \mathbf{c} , define a communication network \mathcal{N}_c .

As a further refinement, a *communication efficiency network* for a given context \mathbf{c} is defined to be a communication network for which the link weights are communication efficiencies; i.e., $w = \varphi_c$. Because $\varphi_c(a, b)$ is undefined if $s_c(a) = 0$, a further restriction is imposed on communication efficiency networks that if $s_c(a) = 0$ then $(a, b) \notin E$, for any $b \in V$.

2.3 Network efficiencies

In the domains under consideration, messages are broadcast by each agent to all other agents. Consequently, a communication efficiency network includes all possible links between nodes (except those which were excluded above: i.e., links for which the communication efficiency is undefined and links which are reflexive), that is $E = \{(a, b) \in V \times V \mid s_c(a) \neq 0 \ \& \ a \neq b\}$.

For each communication network, let the total number of messages sent by all agents in given context \mathbf{c} during a given period be denoted S_c ; i.e., $S_c = \sum_{a \in V} s_c(a)$. Similarly, let the total number of messages received by all agents in the same context over the same period be: $R_c = \sum_{(a,b) \in V \times V} r_c(a, b)$.

In general, each message sent by some agent will be received by some subset of the $|V| - 1$ other agents. If all messages broadcast were received successfully by all other agents, then the total number of messages received is simply the product of the total number of messages sent with the total number of agents receiving those messages, that is

$$R_c = (|V| - 1) \times S_c.$$

Table 1 Roles matched with specific player numbers

Role	Assigned number(s)
Goalkeeper	{1}
Defence	{2, 3, 4, 5}
Mid-field	{6, 7, 8}
Attack	{9, 10, 11}

If some messages are lost, then there will be fewer total messages received: i.e., $R_c \leq (|V| - 1) \times S_c$; or equivalently

$$\frac{R_c}{(|V| - 1) \times S_c} \leq 1.$$

In general, this would imply a reduced network communication efficiency.

Consequently, the *overall efficiency*, Φ , of a communication network \mathcal{N} is defined as the ratio of all messages sent to messages received per agent/node.¹ In particular, given a context \mathbf{c} :

$$\Phi(\mathcal{N}_c) = \frac{R_c}{(|V| - 1) \times S_c}. \tag{2}$$

Observe that for any communications network \mathcal{N} , it follows that $0 \leq \Phi(\mathcal{N}) \leq 1$, and that $\Phi(\mathcal{N}) = 1$ iff $R = (|V| - 1) \times S$.

2.4 Application domain: football simulation

The specific application domain for this project was the RoboCup 2D simulator [17]. The agents were players (i.e., $|\mathbf{A}| = 11$) and the period was the duration of a match, comprising 6000 discrete match cycles. In Table 1, we specify the player role matched with their number.

Attention was restricted to three message types:

- pass messages (p) which contain the intended receiver’s player number and the pass’s destination;
- ball messages (B) which contain information about the position and velocity of the ball; and
- three-player messages (R) which contain information about three players (teammates or opponents).

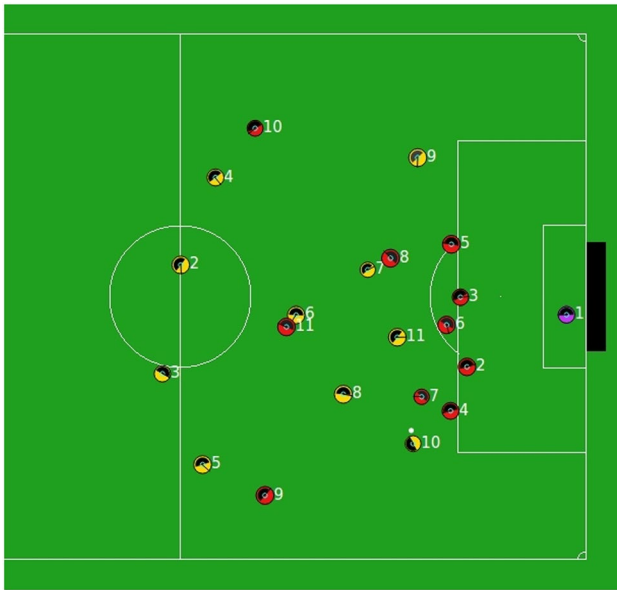
Situations are classified according to two criteria:

- ball possession status: our team in possession of the ball (BPT) or opposition team in possession (BPO);
- field location status: ball in our team’s front half (FH) or ball in our back half (BH).

¹ Excluding the message’s sender.

Table 2 Domain classes and elements

Class	Description	Elements
A	Players	1, 2, . . . , 11
M	Message types	p, B, R
S	Situation types	BPT-FH, BPT-BH, BPO-FH, BPO-BH

**Fig. 1** Match situation

The domain-specific elements for the simulator are summarised in Table 2.

For example, a match situation in which the yellow team (via player yellow #10 at the bottom-right) is in possession of the ball in their front half (i.e., a situation of type BPT-FH) is shown in Fig. 1.

The contexts in C are comprised of combinations of message and situation types: e.g., p-BPT-FH, R-BPO-BH, etc.

2.5 Data generation

Each match played on the simulator generates logs of all player actions and events, including communication primitives (*say* and *hear*). The agent/player code is *agent2d*: the well-known base code developed by Akiyama et al. [5], slightly modified to generate additional data about each match situation: possession (categories BPT and BPO) and field location (categories FH and BH). These data were used to calculate s_c and r_c values, and ultimately communication efficiencies for each pair of agents as determined by (1). For each match communication,

efficiency networks were constructed for each context, and the overall network efficiency was computed according to (2). We stress that in practice communication efficiencies generated through *agent2d* are not functionally predetermined by specific contexts, each of which describes a broad class of instances, varying in terms of the urgency of communications. Furthermore, communication efficiency may significantly differ for any given context due to inherent sensor and actuation noise, fragmentation of simulated sensor fields, and incompleteness of available data.

2.6 Performance measures

The team being analysed was initially matched against an identical opponent. Due to the stochastic nature of the simulator's match model, variability in the match outcomes across matches was observed.

To assess overall team performance, several *team performance measures* were recorded for each match:

- *goals scored*: the number of goals scored by the team during a match;
- *goals conceded*: the number of goals conceded by the team during a match; and
- *goal difference*: the difference between goals scored and goals conceded.

One thousand (1000) matches were played and goals scored and goal conceded for the team were recorded, from which the goal difference was derived. For the remainder of this work, we interchangeably refer to these three performance measures as:

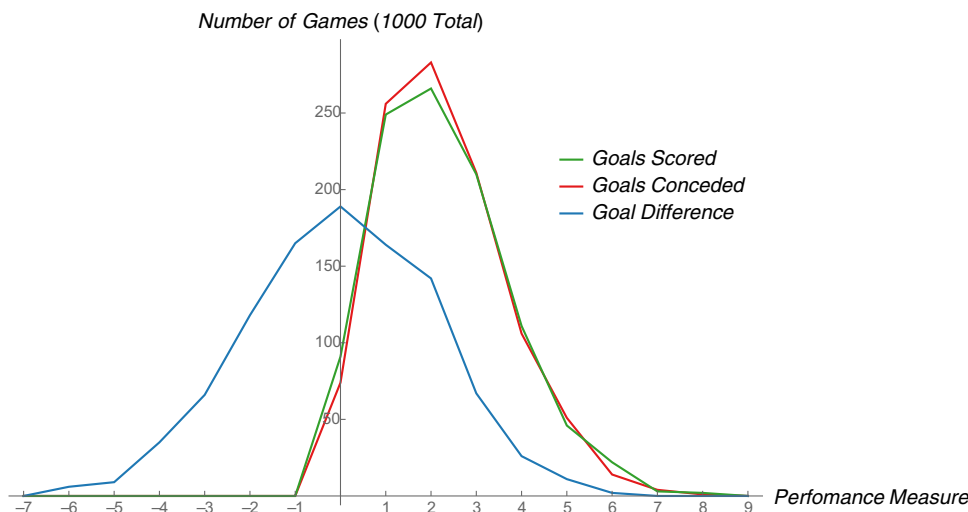
$$\begin{aligned} \mathcal{P}_1 &\equiv \text{goals scored}, & \mathcal{P}_2 &\equiv \text{goals conceded}, \\ \mathcal{P}_3 &\equiv \text{goal difference}. \end{aligned} \quad (3)$$

In Fig. 2, we give the densities of the three performance measures for the 1000 simulation runs. As both simulated teams are equally matched, we find no surprise that both the goals scored and goals conceded densities are very similar, and the goal difference density is approximately normal.

2.7 Regression analysis of the data

Our main results and corresponding narrative are derived from statistical exploration of the data derived from the baseline team settings. Specifically, for the 1000 games simulated, we perform multiple linear regression [45], attempting to explain the independent performance measures of *goals scored/goals conceded/goal difference*, using a simple linear relationship involving the communication efficiencies $\varphi_c(a, b)$. Thus, for each of the four

Fig. 2 Densities of the performance measures for the 1000 simulated games



possible situations listed in Table 2, we construct the following linear relationship:

$$P_i = \sum_{m \in M} \sum_{\substack{a, b \in \mathcal{BIC} \\ a \neq b}} \beta_{a,b}^{(m)} \varphi_c^{(m)}(a, b), \quad i \in \{1, 2, 3\} \quad (4)$$

which maximises the correlation with P_i . \mathcal{BIC} in the summation over players a and b indicates that candidate connections used to build the linear model are chosen using Bayesian Selection Criterion (BIC) [67]. Although many have taken the position that the alternative Akaike Information Criterion (AIC) [4] is superior in model selection to BIC (refer to [2, 79] for examples), we found that, although applying AIC reported higher correlations, a significant proportion of communication efficiencies had coefficient estimates β which were not statistically significantly different from zero. Hence, we opted to apply BIC for model selection.

To the best of our knowledge, the method of measuring inter-agent communication efficiencies over multiple simulation runs and applying linear regression to correlate with performance measures to expose the statistically meaningful links is new in the literature. The method is relatively simple and intuitive to grasp, with the ability to be performed with most simulation engines as although we specifically use *agent2D* to generate the data, we effectively treat it as a ‘black box’. The actual communication protocol of *agent2D* typically rotates the sender agent across all 11 team players over 11 cycles interval (in a predefined order synchronised by the current game cycle number). However, some players may choose to communicate ‘out of turn’ in exceptional circumstances when the available data warrants some urgency, e.g., a player is near the off-side line. In addition, each player may choose to receive messages only from a specific sender, by temporarily

setting its *attention* variable, again under specific conditions. All such exceptions depend on a set of pre-programmed conditions that may be met at a given cycle, and varying such conditions, by modifying the corresponding numerical threshold parameters, would result in changing the communication protocol, modifying the ‘black box’. We will briefly discuss such modifications, each constituting a separate design point, in our conclusions.

It is our hope that other agent-based model research teams, whether in RoboCup or in more general applications, see the utility in our method as a means of exposing important and possibly counterintuitive communications which correlate with performance. It is important to note, however, that although we do report correlation values, we are not fixated on obtaining the maximal value of correlation to obtain a predictive linear model. Rather, our aim is to establish the underlying meaning and motifs behind statistically significant communication links.

3 Results and discussion

We perform multiple linear regression on the communication data for each of the situations in Table 2. However, we have disallowed any communication *received* by the goalkeeper, i.e., $b \neq 1$ in Eq. (4). This is due to statistical artefacts dominating correlation scores *after* the scoring of goals.

Before we present our findings of the regression analysis in their entirety, we shall begin with the situation BPO-BH, concentrating on the performance measure of *goal difference* presented in Table 3 and discuss some necessary subtleties.

The *t*-ratio in Table 3 is the estimate of the coefficient β , divided by the standard error associated with the estimated coefficient (for more information, refer to [45, 76]).

Table 3 Results from performing multiple linear regression (applying BIC) with the independent performance measure of goal difference (\mathcal{P}_3) in the situation BPO-BH

m	a	b	$\beta_{a,b}^{(m)}$	Std. error	t -ratio
B	11	6	8.04	1.52	5.30
B	3	5	-5.39	1.16	-4.65
R	11	4	5.48	1.43	3.84
R	11	2	5.16	1.46	3.53
R	7	8	4.17	1.20	3.46
B	2	5	-8.45	2.67	-3.16
R	10	4	3.65	1.17	3.13
R	1	3	-3.20	1.09	-2.94
B	11	5	4.96	1.85	2.68
B	1	5	-3.36	1.26	-2.68
B	2	3	-6.12	2.39	-2.56
B	4	10	1.95	1.09	1.79

Importantly, the t -ratio determines the statistical significance of each coefficient estimate, and an absolute value greater than 1.96 equates to the coefficient being different from zero with greater than 95% confidence. As explained in Sect. 2.7, we opted to apply BIC instead of AIC for model construction mainly due to the fact that the application of AIC leads to a significant proportion of coefficient estimates β having t -ratios less than 1.96—hence not statistically significantly different from zero. In Table 3, we can see that only one coefficient estimate has such a t -ratio.

In addition, it is important to highlight the difference between *positive* and *negative* coefficient estimates in Table 3. Positive estimates mean that the linear model predicts an increase/decrease in performance if the communication channel corresponding to the estimated

coefficient increases/decreases efficiency. Conversely, negative estimates mean that the linear model predicts a decrease/increase in performance if the communication channel corresponding to the estimated coefficient increases/decreases efficiency.

3.1 Network representations: individual situations

We now present the full results of the regression analysis in the form of network diagrams. Starting with the situation BPO-BH in Fig. 3, we note that communications with positive coefficient estimates are shown as solid lines, and negative coefficient estimates are shown as dashed lines. In addition, interpreting the *goals conceded* graph needs care; solid edges (positive coefficient estimates) correlate to *goals being conceded*, which is a traditionally counter intuitive way of thinking about performance. We also note that for each graph, we give the R^2 (the Pearson correlation squared) value, which measures the percentage of variation in the data captured by the linear model.

Focusing on *goals scored* as the performance measure in Fig. 3, we see that communications with positive correlations are almost exclusively with agents over a long-distance (e.g. $11 \rightarrow 4, 10 \rightarrow 5$), indicating that useful data are being transferred to recipients which is outside of their field of view. The obvious counter example to this is player 11, the centre-forward, communicating with mid-fielders 6 and 8. For *goals conceded*, only connections with negative coefficient estimates are present. In addition, there is visually no correlation between *goals conceded* and *goal difference* for this situation, i.e., no links in one graph are present in the other. However, there is quite good agreement between *goals scored* and *goal difference*.




Fig. 3 Graphical representation of performing regression in the situation BPO-BH for all performance measures. Note that *goal difference* corresponds to the data in Table 3 and *thin edges* represent statistically insignificant links




Fig. 4 Graphical representation of performing regression in the situation BPT-BH for all performance measures. Note that *thin edges* represent statistically insignificant links




Fig. 5 Graphical representation of performing regression in the situation BPO-FH for all performance measures

We also notice that the majority of communications from defenders and the goalie are negatively correlated with performance, possibly due to the received information quickly becoming obsolete. The converse statement is also true: that the majority of communications from mid-fielders and forwards are positively correlated with performance. This phenomenon is likely due to agents reporting changes in ball and player positions (likely during opponent’s passes) and hence enabling better ball interceptions and counter-attacks.

Finally, in *goals conceded*, we notice a two-hop motif (7 → 11 → 3, 7 → 11 → 2, and 6 → 8 → 5) which possibly improves the quality of data, about the ball and other players, coming to the defenders.

For the situation BPT-BH in Fig. 4, we now see the appearance of message type p, associated with passing, in the networks. In addition, we also see that the graph for *goal difference* contains a proportion of links which appear in both *goals scored* (3 → 10 p) and *goals*

conceded (9 → 10 R for instance), unlike the previous situation. We also see that *goals conceded* has the highest R^2 and is the only performance measure which does not contain disconnected clusters.

In *goals scored*, there is a motif of long-distance communications, 1 → 11, 4 → 8, and 3 → 10, carrying data about the ball and players which are likely to be outside of the recipient’s field of view, and helping to build-up a counter-attack. For negatively correlated B communications in *goals conceded*, the short-range chain motifs (9 → 11, 4 → 7 → 6 → 8 → 3 → 5 → 6 and 4 → 7 → 8 → 3 → 5 → 6) help teammates to keep possession of the ball. A general trend across all performance measures for this situation is that communications sent by forwards are not helpful as they possibly conflate data about ball possession.

Focusing on Fig. 5 which presents the situation BPO-FH, we note that this example has some of the lowest

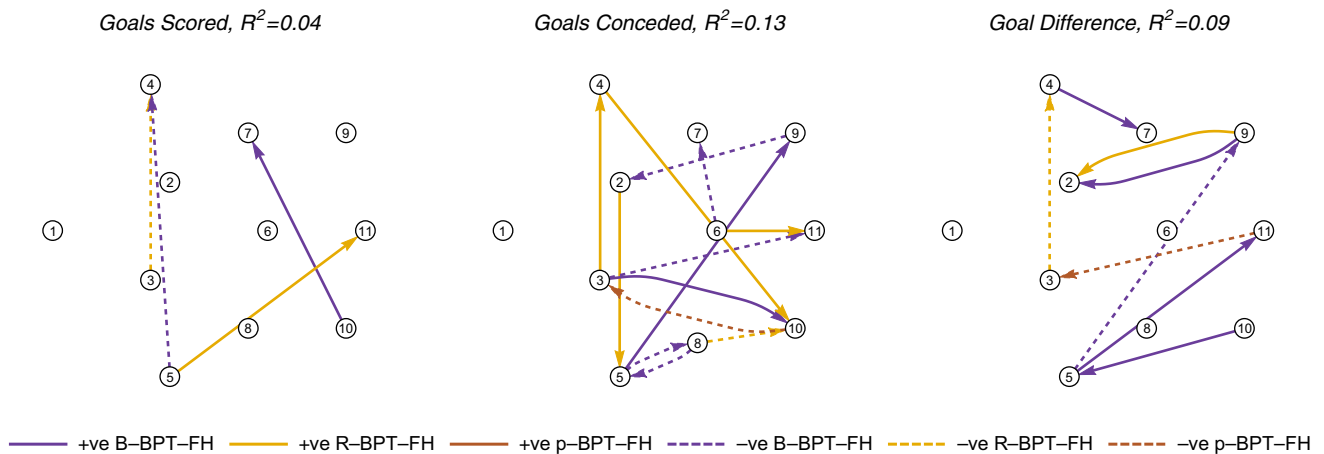


Fig. 6 Graphical representation of performing regression in the situation BPT-FH for all performance measures

Pearson correlation values. Specifically, we notice that the BIC algorithm has failed to detect any meaningful correlation between communication efficiencies between agents and *goals conceded*. From the other two performance measures, we can see that, generally, communication is not beneficial for this situation.

In Fig. 6, we present the situation BPT-FH. Focusing on *goals scored*, we can see just how little communication efficiency seems to correlate with this performance measure. This result is not entirely surprising given the simplicity of the scoring tactics used by the *agent2d* team (employing the 4-3-3 formation, with three mid-fielders, and three forwards, and being dominated by side crosses). Specifically, however, the positive $5 \rightarrow 11$ link for this measure may help the centre-forward stay *onside*.

For *goals conceded* we see that there are some beneficial communication motifs to defenders and mid-fielders ($8 \rightarrow 5$, $5 \rightarrow 8$, $10 \rightarrow 3$, $6 \rightarrow 7$, and $9 \rightarrow 2$) which improve data quality and *may* help to prevent opponents' counter-attacks. The beneficial links $8 \rightarrow 10$ and $3 \rightarrow 11$ for *goals conceded* also suggest that forward agents are contributing to preventing counter-attacks from the opposing team. Correspondingly, the mid-range communications amongst defenders ($2 \rightarrow 5$ and $3 \rightarrow 4$) in *goals conceded* are negatively correlated with performance, indicating that they are potentially enabling poor quality passes via propagation of out-dated players' positions.

Finally, focusing on *goal difference*, there is a general trend that lateral communications across the field (e.g. $3 \rightarrow 4$, $5 \rightarrow 9$) are negatively correlated with performance, and longitudinal communications up the field (e.g. $9 \rightarrow 2$, $4 \rightarrow 7$) are positively correlated.

In general, BPO-BH has highest R^2 for *goals scored*, this may seem surprising due to this situation being the

most defensive. However, a sizable proportion of goals are scored as a consequence of a contest between two players won by the defending team; the team's forwards playing on the sides (wing-forwards) very quickly position themselves to receive the ball, and progress to make a cross resulting in scoring a goal. If we contrast this result with *goals scored* in BPT-FH, which has negligible R^2 , and few network links, we can conclude that communication contributes very little to goal-scoring while in an attacking situation.

In addition, we see that BPT-BH has the highest R^2 for *goals conceded*. This shows that, when the players are in relatively defensive situations, maintaining possession is crucial to prevent the goals scored by the opposing team. Contrasting this with *goals conceded* in BPO-BH, the most defensive situation, which has a small R^2 and few connections, we can conclude that communication contributes very little to conceding goals in the most defensive situation.

3.2 Network representations: all situations

We now perform regression analysis on *all* the communication data, regardless of situation, presented graphically in Fig. 7. Due to the larger amount of links that are being accepted by the BIC algorithm, it is necessary to split the links with +ve and -ve coefficient estimates in separate networks.

Focusing on *goals scored*, for the top-left graph, over half of the links are derived from the situation BPO-BH, with twice as many message type R links as there are B links. Interestingly, we see no appearance of message type p in this graph. The corresponding graph in the bottom-left, giving -ve coefficient estimates also has the majority of its links

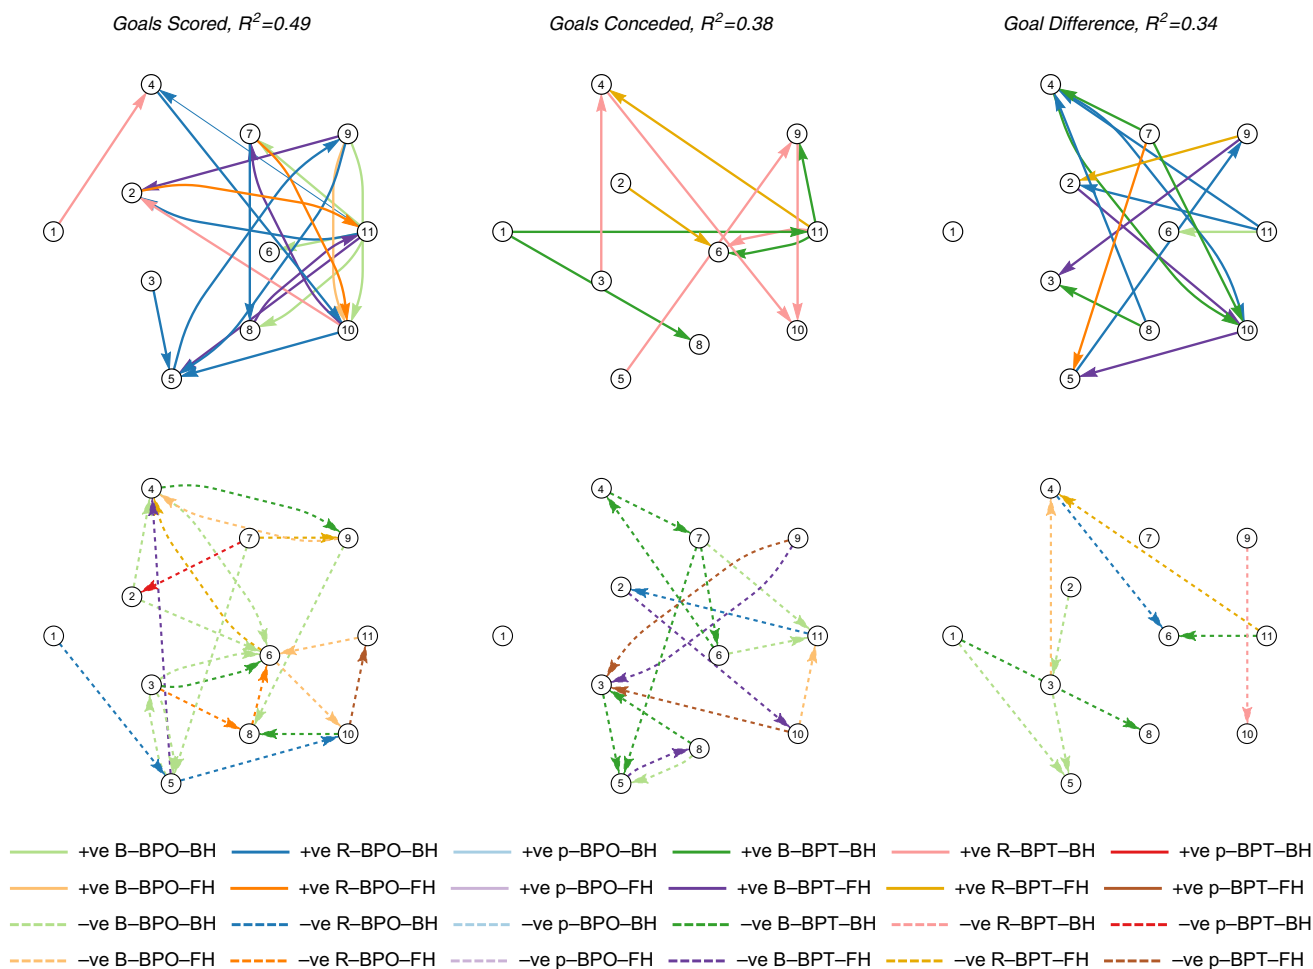


Fig. 7 Graphical representation of performing regression for all situations and performance measures. Note that +ve and -ve coefficient estimates have been segregated for clarity and *thin edges* represent statistically insignificant links

derived from the situation BPO-BH. This time, we have four times as many message type B links as there are R links for this situation—the reverse of the corresponding +ve estimate graph. This time, however, we have two links (7 → 2 and 10 → 11) for message type p implying that communications about passing are largely detrimental to scoring.

For *goals conceded*, the top-middle graph only contains three types of links: B-BPT-BH, R-BPT-BH, and R-BPT-FH. The -ve coefficient estimate graph in the bottom-middle is much more complicated; nevertheless, looking at both graphs together, we can conclude that more R-type communication whilst in possession of the ball (BPT) correlates with more goals being conceded. Looking at the corresponding -ve estimate graph, we can make the equivalent claim that less type B communications in situations BPT-BH and BPT-FH correlate to less goals being conceded. This is due to the graph consisting of over 80% of these (B) types of communication links.

Finally, for *goal difference*, we can see that when the opponent has possession of the ball (BPO) in the top-right graph, all but one of the links correlating with performance are of message type R. Likewise, when in possession of the ball (BPT), all but one of the links are of message type B. These two results distinguish the importance of these two message types based on who has ball possession. For the corresponding -ve coefficient estimate graph on the bottom-right, we can see that the majority of links are derived from the back half (BH) and are of message type B, indicating the communication about the ball while in the back half negatively correlates with performance.

For the majority of the graphs in Fig. 7, we largely see that communications from the defenders is not beneficial. In addition, communications to the centre mid-field (player 6) are not beneficial. Nevertheless, communications to the forwards are generally useful for scoring, as are the

Table 4 Results from performing regression (applying BIC) with the independent performance measure of *goal difference* (\mathcal{P}_3) in the situation BPO-BH on aggregated data

m	Sender	Receiver	$\beta^{(m)}$	Std. error
B	Attack	Mid-field	0.26	0.04
R	Goalie	Defence	-0.29	0.05
R	Mid-field	Defence	0.19	0.04
R	Attack	Defence	0.17	0.04
R	Goalie	Attack	0.18	0.06

communications to the defenders from the mid-fielders and forwards.

3.3 Aggregation of players roles: principal components

In this section, we again give network representations of multiple linear regression analysis, but instead of considering communications between individual players, we first aggregate players in terms of their role (defenders: Defence, mid-fielders: Mid-field, and forwards: Attack). For each message M and situation S type, we group both sender and receiver in terms of player roles, as shown in Table 1. As an example, the grouping for $\{\text{sender, receiver}\} = \{\text{Defence, Mid-field}\}$ is the subset of communication efficiencies, $\varphi_c^{(m)}(a, b)$ with $a \in \{2, 3, 4, 5\}$, $b \in \{6, 7, 8\}$. Thus, for each specific message and situation type, we obtain 12 different subsets of sender-receiver groupings (recall that we discount the goalkeeper as a message receiver due to statistical artefacts).

We then reduce the dimension of each aforementioned grouping to one through performing principal components analysis (PCA) [36] to each grouping, and only keeping the component which contains the most variance (the principal component). Through performing dimensional reduction in this way,² we obtain an aggregated communication efficiency which is most representative of the original data as it has the highest possible variance. Thus, for each message and situation type, we reduce our original data down to 12 aggregated communication efficiencies.

Finally, for each situation type, we perform regression on this aggregated data against each of the performance measures \mathcal{P}_i , $i \in \{1, 2, 3\}$. We give an example in Table 4 for the situation BPO-BH and performance measure of

goal difference, which is the aggregated equivalent to what was presented in Table 3.

Starting with the situation BPO-BH in Fig. 8, we note the appearance of self-loops for the first time. Focusing on *goals scored*, we see that type R communications amongst forwards negatively correlate with performance. Interestingly, referring back to the corresponding graph in Fig. 3, there were no -ve coefficient estimates amongst the forward grouping, indicating that this new phenomenon is a result of the PCA procedure combined with the BIC algorithm. Nevertheless, the self-loop amongst the mid-fielders in the *goals conceded* graph is visible in the corresponding graph in Fig. 3.

For *goal difference*, although it visibly correlates quite well with *goals scored* (containing no links from *goals conceded*), the links Goalie \rightarrow Attack and Mid-field \rightarrow Defence are not present in the corresponding graph in Fig. 3. This indicates that even though the individual player-to-player communication links are not significant, the role-to-role aggregation provided by PCA shows these links in a new light and exposes their worth.

For the situation BPT-BH in Fig. 9, we see that message type p does not feature, which is a departure from Fig. 4 when performing regression on individual players. We also see that the graph for *goal difference* contains some links which appear in both *goals scored* (Attack \rightarrow Defence) and *goals conceded* (Mid-field \rightarrow Mid-field for instance), unlike the previous situation in Fig. 8. Similar to the previous situation, however, certain links are present in all three graphs (Goalie \rightarrow Defence in *goals scored* for instance) which are not present in corresponding graphs in Fig. 4: a result of the PCA procedure. Thus, although individual links from the goalkeeper to the defenders are deemed statistically insignificant, the sum of these links through the PCA aggregation procedure now appears as significant. We do not show the results for the situations BPO-FH and BPT-FH as the correlations are small and the graphs are quite sparse.

Finally, in Fig. 10, we give the graphs obtained through performing regression on all the aggregated communication data, regardless of situation. In general, we see that these networks paint a similar picture to what was noted in Fig. 7. Focusing on *goals scored*, most of the links are from situation BPO-BH. For the +ve coefficient estimates, most are again of message type R. For the -ve coefficient estimates, however, we now have a more even mix of message types B and R.

For *goals conceded*, we see that there are only messages B-BPT-BH and R-BPT-BH correlating positively with performance. For the -ve coefficient estimates, all but one of the messages relate to having ball possession (BPT). In addition, as there are no type R-BPT messages negatively correlated with performance, we can draw the same

² One alternative procedure which would maximise correlation would be to perform regression to each grouping against the performance measures \mathcal{P}_i .

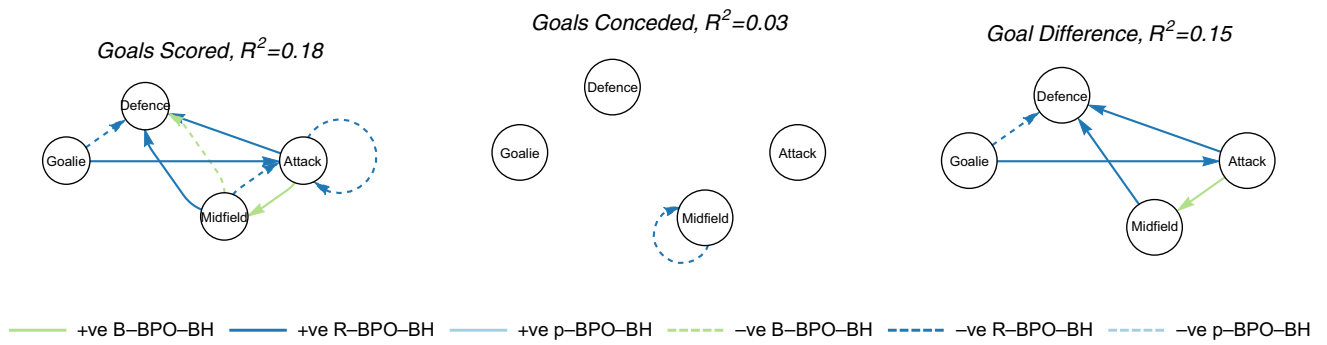


Fig. 8 Graphical representation of performing aggregated regression in the situation BPO-BH for all performance measures. Note that *goal difference* corresponds to the data in Table 4

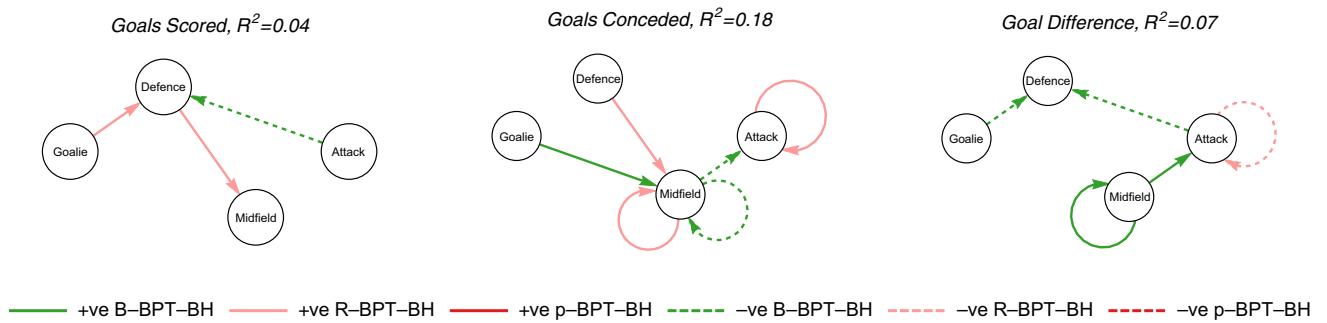


Fig. 9 Graphical representation of performing aggregated regression in the situation BPT-BH for all performance measures




Fig. 10 Graphical representation of performing aggregate regression for all situations and performance measures

conclusion as the corresponding graph in Fig. 7; that more type R communication whilst in possession of the ball correlates to more goals being conceded.

Finally, upon visual inspection of the graph presented in *goal difference*, we can see that it correlates very well

to the difference of the previous two performance measure graphs. The only departure to this being Goalie → Attack -ve B-BPT-BH (which does not appear in the corresponding graph in Fig. 7) and Attack → Defence +ve R-BPT-FH (which does appear in Fig. 7).

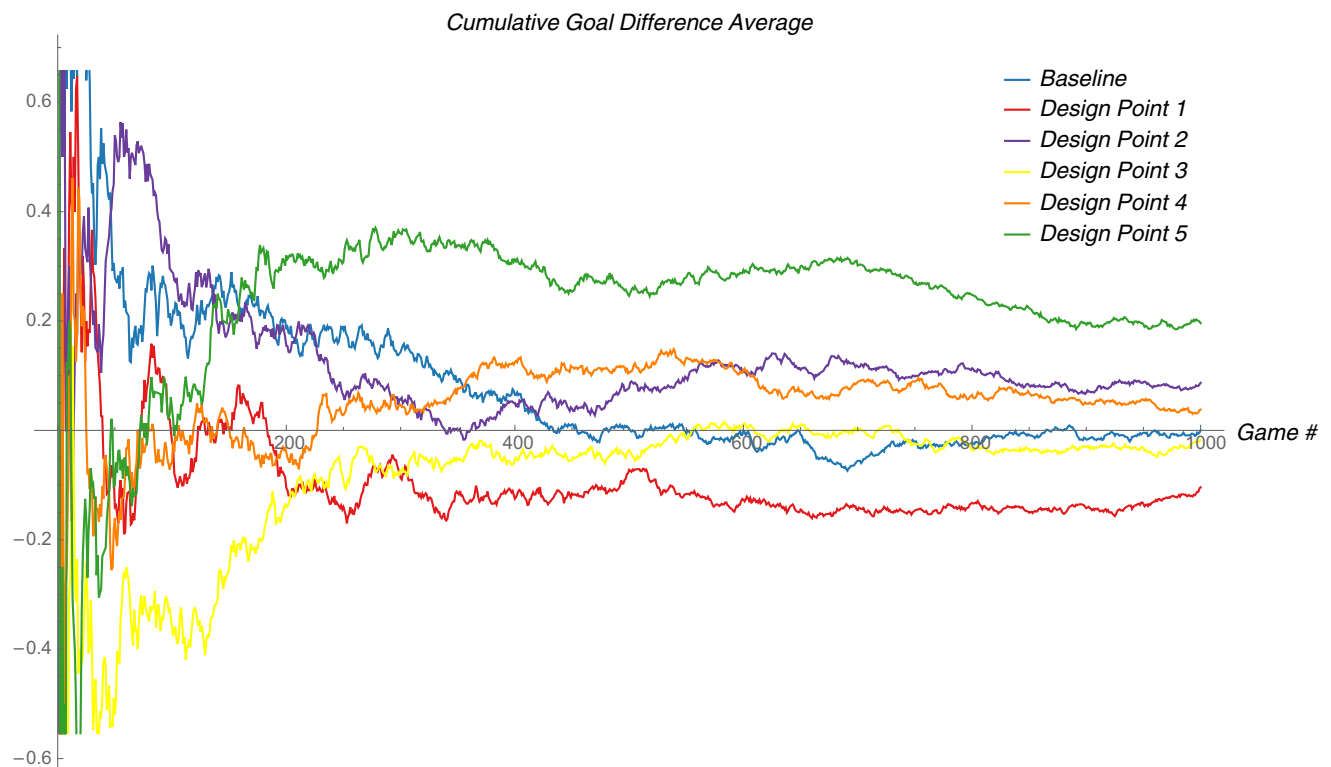


Fig. 11 Cumulative average of *goal difference* for six different design points over 1000 simulated games. Note that the *baseline* design point (*blue*) is the one explored in this work

4 Conclusions and future work

In this exploratory study, we measured the inter-agent communication efficiencies and performance measures from multiple runs of RoboCup Soccer Simulation 2D League. From these data, we generated agent communication networks, for different situations *S* and message types *M*, through performing multiple linear regression of the communication data against the performance measures. These *functional* networks enabled us to determine the player-to-player communications which correlated (positively and negatively) the strongest with the corresponding performance measure. Visual inspection of these networks revealed relevant motifs which highlight player tactics contributing to performance of a *structured* team comprising multiple agents.

Thus, the study analysed both the *structural* and *functional* connectivity, and offered a method for deriving functional communication networks, for various situation types and structured agent roles, based on the underlying communication data correlated with the overall performance. In other words, the functionality is defined directly in terms of overall team performance, and so, the topology of the resultant functional networks is dependent on the global outcomes. One immediate utility of our technique is its ability to expose counterintuitive motifs; for example, during BPO-BH,

communication between defenders and the goalie is negatively correlated with performance, whereas communications between mid-fielders and forwards are positively correlated.

Furthermore, through utilising PCA on the data to aggregate the players into roles, the corresponding networks enabled us to determine the macroscopic role-to-role communications which correlated with performance. Though largely reinforcing the narrative generated by the player-to-player networks, visual inspection of the macroscopic networks also revealed some interesting communications, mostly with the goalkeeper, which only appear due to the aggregation of players into roles. This shows that although individual player links may be deemed statistically insignificant on their own, the true worth of the links only appears when roles are aggregated through PCA, or some equivalent technique.

It is our hope that the analysis presented in this work, which attempts to expose the non-trivial links between communication and performance in a multi-agent-based setting, will be used as a platform for future studies in this area. Specifically, the communication protocol and its thresholds are controlled in *agent2D* by numerical global input parameters. An obvious follow-up study would include generating, and comparing, communication data for simulation runs through a systematic variation of these

global input parameters (akin to a *data-farming* experiment, where particular input parameter combinations are referred to as *design points* [20]). In Fig. 11, we offer the cumulative average of *goal difference* for six different design points, where we have varied the global input parameters of the home team, which changes various communication policies and protocols. We note that the *baseline* design point was focused in this work. A systematic data-farming experiment would allow an optimisation of these global communication variables against performance measures.

This would be an obvious avenue of generating more effective communication policies and protocols amongst the players. Indeed, it is our intent that the methodology detailed in this work be applied by various RoboCup teams as a means of highlighting important communication links, and generating more effective inter-agent coordination. It would also be of interest to perform similar analysis on other adversarial multi-agent-based simulators where communication is paramount to performance [33]. In such a study, with potentially many more agents under consideration, the corresponding networks would be more complicated, thus requiring more sophisticated forms of analysis (such as understanding the narrative generated from social network analysis metrics [80]) to expose the communication motifs. In addition, in scenarios involving significantly more agents, performing PCA to expose macroscopic communication links between relevant functions may yield more utility than seen in this work as they would enable the possibility of effective visual inspection.

Finally, although our technique performed multiple linear regression on the communication efficiencies against performance, most information-theoretic measures typically use some variants of entropy and hence are nonlinear. Another follow-up study would involve determining some nonlinear transform of the communication data which meaningfully increased the Pearson correlation values with performance measures.

Acknowledgements The authors wish to thank Alexander Kalloniatis for stimulating discussions. This work was funded by Defence Science and Technology Group's Trusted Autonomous Systems Strategic Research Initiative (Project Tyche).

References

1. Abreu PH, Moura J, Silva DC, Reis LP, Garganta J (2012) Performance analysis in soccer: a Cartesian coordinates based approach using RoboCup data. *Soft Comput* 16(1):47–61
2. Aho K, Derryberry D, Peterson T (2014) Model selection for ecologists: the worldviews of aic and bic. *Ecology* 95(3):631–636
3. Ajitha S, Suresh Kumar TV, Rajanikanth K (2013) A quantitative framework for early prediction of cooperation in multi-agent systems. *ICTACT J Soft Comput* 3(4):587–595
4. Akaike H (1974) A new look at the statistical model identification. *IEEE Trans Autom Control* 19(6):716–723
5. Akiyama H (2010) Agent2D Base Code. www.rctools.sourceforge.jp
6. Akiyama H, Nakashima T (2013) HELIOS2012: RoboCup 2012 Soccer Simulation 2D League Champion. In: Chen X, Stone P, Sucar L, van der Zant T (eds) *RoboCup 2012: Robot Soccer World Cup XVI*. Springer, Berlin, pp 13–19
7. Bai A, Wu F, Chen X (2015) Online planning for large markov decision processes with hierarchical decomposition. *ACM Trans Intell Syst Technol* 6(4):45
8. Bai A, Chen X, MacAlpine P, Urieli D, Barrett, Stone P (2012) Wrighteagle and ut austin villa: RoboCup 2011 simulation league champions. In: Röfer T, Mayer N, Savage J, Saranlı U (eds) *RoboCup 2011: Robot Soccer World Cup XV*. Lecture Notes in Artificial Intelligence. Springer, Berlin, pp 1–12
9. Becker R, Carlin A, Lesser V, Zilberstein S (2009) Analyzing myopic approaches for multi-agent communication. *Comput Intell* 25(1):31–50
10. Bernstein D, Zilberstein S, Immerman N (2002) The complexity of decentralized control of markov decision processes. *Math Oper Res* 27(4):819–840
11. Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU (2006) Complex networks: structure and dynamics. *Phys Rep* 424(4):175–308
12. Budden D, Prokopenko M (2013) Improved particle filtering for pseudo-uniform belief distributions in robot localisation. In: Behnke S, Veloso MM, Visser A, Xiong R (eds) *RoboCup 2013: Robot Soccer World Cup XVII*, vol 8371. Lecture notes in artificial intelligence. Springer, Berlin
13. Budden D, Wang P, Obst O, Prokopenko M (2014) Simulation leagues: analysis of competition formats. In: Bianchi RAC, Akin HL, Ramamoorthy S, Sugiura K (eds) *RoboCup 2014: Robot Soccer World Cup XVIII*, vol 8992. Lecture notes in computer science. Springer, Berlin, pp 183–194
14. Budden D, Wang P, Obst O, Prokopenko M (2015) RoboCup simulation leagues: enabling replicable and robust investigation of complex robotic systems. *IEEE Robot Autom Mag* 22(3):140–146
15. Butler M, Prokopenko M, Howard T (2001) Flexible synchronisation within RoboCup environment: a comparative analysis. In: Stone P, Balch T, Kraetzschmar G (eds) *RoboCup 2000: Robot Soccer World Cup IV*, vol 2019. Lecture notes in artificial intelligence. Springer, Berlin, pp 119–128
16. Candea C, Hu H, Iocchi L, Nardi D, Piaggio M (2001) Coordination in multi-agent robocup teams. *Robot Auton Syst* 36(2):67–86
17. Chen M, Foroughi E, Heintz F, Huang Z, Kapetanakis S, Kostiadis K, Kummeneje J, Noda I, Obst O, Riley P, Steffens T, Wang Y, Yin X (2001) *RoboCup Soccer Server*. <http://wwf.cs.virginia.edu/documentation/manual.pdf>
18. Chickering DM (2002) Learning equivalence classes of Bayesian-network structures. *J Mach Learn Res* 2(Feb):445–498
19. Chou WYJ, Marsh L, Gossink D (2009) Multi-agent coordination and optimisation in the RoboCupRescue project. In: 18th World IMACS/MODSIM Congress. Cairns, Australia (July) (Citeseer)
20. Cioppa TM, Lucas TW (2007) Efficient nearly orthogonal and space-filling latin hypercubes. *Technometrics* 49(1):45–55
21. Cliff O, Prokopenko M, Fitch R (2016) An information criterion for inferring coupling of distributed dynamical systems. *Front Robot AI* 3:71
22. Cliff OM, Lizier JT, Wang P, Wang XR, Obst O, Prokopenko M (2017) Quantifying long-range interactions and coherent structure in multi-agent dynamics. *Artif Life* 23(1):34–57
23. Cliff OM, Lizier JT, Wang XR, Wang P, Obst O, Prokopenko M (2014) Towards quantifying interaction networks in a

- football match. In: Behnke S, Veloso M, Visser A, Xiong R (eds) *RoboCup 2013: Robot World Cup XVII*, vol 8371. Lecture notes in computer science. Springer, Berlin, pp 1–12
24. Fewell JH, Armbruster D, Ingraham J, Petersen A, Waters JS (2012) Basketball teams as strategic networks. *PLoS One* 7(11):e47445
 25. Frias-Martinez V, Marcinkiewicz M, Parsons S, Sklar E (2004) Using multiagent coordination techniques in the robocup four-legged league. In: *Proceedings of the AAAI spring symposium on bridging the multi-agent and multi-robotic research gap*
 26. Gan SK, Fitch R, Sukkariyah S (2014) Online decentralized information gathering with spatial-temporal constraints. *Auton Robot* 37(1):1–25
 27. Ghahramani Z (1998) Learning dynamic bayesian networks. In: Giles C, Gori M (eds) *Adaptive processing of sequences and data structures*. Springer, Berlin, pp 168–197
 28. Gutiérrez C, García-Magariño I (2009) A metrics suite for the communication of multi-agent systems. *J Phys Agents* 3(2):7–14
 29. Gutiérrez C, García-Magariño I, Fuentes-Fernández R (2011) Detection of undesirable communication patterns in multi-agent systems. *Eng Appl Artif Intell* 24(1):103–116
 30. Haker M, Meyer A, Polani D, Martinecz T (2002) A method for incorporation of new evidence to improve world state estimation. In: Birk A, Coradeschi S, Tadokoro S (eds) *RoboCup 2001: Robot Soccer World Cup V*, vol 2377. Lecture notes in computer science. Springer, Berlin, pp 362–367
 31. Hausknecht M, Mupparaju P, Subramanian S, Kalyanakrishnan S, Stone P (2016) Half field offense: an environment for multiagent learning and ad hoc teamwork. In: *AAMAS adaptive learning agents (ALA) workshop*
 32. Howard RA (1966) Information value theory. *Syst Sci Cybern IEEE Trans* 2(1):22–26
 33. Ilachinski A (2004) *Artificial war: multiagent-based simulation of combat*. World Scientific, Singapore
 34. Jennings NR, Sycara K, Wooldridge M (1998) A roadmap of agent research and development. *Auton Agents Multi-agent Syst* 1(1):7–38
 35. Jinyi Y, Ni L, Fan Y, Yunpeng C, Zengqi S (2004) Technical solutions of TsinghuAeolus for robotic soccer. In: Polani D, Browning B, Bonarini A, Yoshida K (eds) *RoboCup 2003: Robot Soccer World Cup VII*. Springer, Berlin, pp 205–213
 36. Jolliffe IT (2002) *Principal component analysis*, 2nd edn. Springer, New York
 37. Kantz H, Schreiber T (2004) *Nonlinear time series analysis*, vol 7. Cambridge University Press, Cambridge
 38. Kitano H, Asada M (1998) The RoboCup humanoid challenge as the millennium challenge for advanced robotics. *Adv Robot* 13(8):723–736
 39. Kitano H, Tambe M, Stone P, Veloso MM, Coradeschi S, Osawa E, Matsubara H, Noda I, Asada M (1998) *The RoboCup synthetic agent challenge 97*. In: Kitano M (ed) *RoboCup-97: Robot Soccer World Cup I*, London, UK. Springer, Berlin, pp 62–73
 40. Kok JR, Spaan MTJ, Vlassis NA (2005) Non-communicative multi-robot coordination in dynamic environments. *Robot Auton Syst* 50(2–3):99–114
 41. Kok JR, Vlassis N, Groen FCA, UvA Trilearn (2003) team description. In: Polani D, Browning B, Bonarini A, Yoshida K (eds) *Proceedings CD RoboCup 2003*, Padua, Italy, July 2003. Springer, Berlin
 42. Lizier J, Heinzle J, Horstmann A, Haynes J, Prokopenko M (2011) Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *J Comput Neurosci* 30(1):85–107
 43. Lizier J, Rubinov M (2012) Multivariate construction of effective computational networks from observational data. *Max Planck Institute for Mathematics in the Sciences* (preprint)
 44. MacAlpine P, Barrett S, Urieli D, Vu V, Stone P (2012) Design and optimization of an omnidirectional humanoid walk: a winning approach at the RoboCup 2011 3D simulation competition. In: *Proceedings of the 26th conference on artificial intelligence, AAAI*
 45. Moore D, McGabe G (1993) *Introduction to the practice of statistics*, 2nd edn. W.H Freedman and Company, New York
 46. Mota L, Reis L, Lau N (2011) Multi-robot coordination using setplays in the middle-size and simulation leagues. *Mechatronics* 21(2):434–444
 47. Nair R, Tambe M, Marsella S (2003) Team formation for reformation in multiagent domains like RoboCupRescue. In: Kaminka G, Lima P, Rojas R (eds) *RoboCup 2002: Robot Soccer World Cup VI. Lecture notes in computer science*, vol. 1395. Springer Verlag, Berlin, pp 150–161
 48. Niehaus C, Röfer T, Laue T (2007) Gait optimization on a humanoid robot using particle swarm optimization. In: *Proceedings of the second workshop on humanoid soccer robots*. IEEE
 49. Noda I, Stone P (2003) The RoboCup soccer server and CMU-nited clients: implemented infrastructure for MAS research. *Auton Agents Multi-agent Syst* 7(1–2):101–120
 50. Panait L, Luke S (2005) Cooperative multi-agent learning: the state of the art. *Auton Agents Multi-agent Syst* 11(3):387–434
 51. Park HJ, Friston K (2013) Structural and functional brain networks: from connections to cognition. *Science* 342(6158):1238411
 52. Peña JL, Touchette H (2012) A network theory analysis of football strategies. In: *Proc Euromech Physics of Sports Conference, 2012* (arXiv preprint). [arXiv:1206.6904](https://arxiv.org/abs/1206.6904)
 53. Prokopenko M, Wang P (2016) Disruptive innovations in RoboCup 2D soccer simulation league: from Cyberoos'98 to Gliders2016. In: Behnke S, Lee D, Sariel S, Sheh R (eds) *RoboCup 2016: Robot Soccer World Cup XX. Lecture notes in artificial intelligence*. Springer, Berlin
 54. Prokopenko M, Wang P, Obst O, Jauregui V (2016) Gliders 2016: Integrating multi-agent approaches to tactical diversity. In: *RoboCup 2016 symposium and competitions: team description papers*. Germany, July, Leipzig, p 2016
 55. Prokopenko M, Ay N, Obst O, Polani D (2010) Phase transitions in least-effort communications. *J Stat Mech Theory Exp* 2010(11):P11025
 56. Prokopenko M, Wang P (2003) Relating the entropy of joint beliefs to multi-agent coordination. In: Kaminka G, Lima P, Rojas R (eds) *RoboCup 2002: Robot Soccer World Cup VI*, vol 2752. Lecture notes in computer science. Springer, Berlin, pp 367–374
 57. Prokopenko M, Wang P (2004) Evaluating team performance at the edge of chaos. In: Polani D, Browning B, Bonarini A, Yoshida K (eds) *RoboCup 2003: Robot Soccer World Cup VII*, vol 3020. Lecture notes in computer science. Springer, Berlin, pp 89–101
 58. Pynadath DV, Tambe M (2002) The communicative multiagent team decision problem: analyzing teamwork theories and models. *J Artif Intell Res* 16:389–423
 59. Rein R, Memmert D (2016) Big data and tactical analysis in elite soccer: future challenges and opportunities for sports science. *SpringerPlus* 5(1):1410–1412
 60. Riedmiller M, Gabel T, Knabe J, Strasdat H (2006) Brainstormers 2D - team description 2005. In: Bredendfeld A, Jacoff A, Noda I, Takahashi Y (eds) *RoboCup 2005: Robot Soccer WorldCup IX. Lecture notes in computer science*, vol. 4020. Springer Verlag, Berlin, pp 219–229
 61. Riedmiller M, Gabel T, Hafner R, Lange S (2009) Reinforcement learning for robot soccer. *Auton Robot* 27(1):55–73
 62. Riley P, Stone P, Veloso M (2001) Layered disclosure: revealing agents' internals. In: Castelfranchi C, Lesperance Y (eds),

- Intelligent Agents VII. Agent Theories, Architectures, and Languages—7th. International Workshop, ATAL-2000, Boston, MA, USA, July 7–9, 2000, Proceedings, Lecture Notes in Artificial Intelligence. Springer, Berlin
63. Roth M, Simmons R, Veloso M (2005) Reasoning about joint beliefs for execution-time communication decisions. In: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems. ACM, pp 786–793
 64. Roth M, Simmons R, Veloso M (2006) What to communicate? Execution-time decision in multi-agent pomdps. Distributed autonomous robotic systems 7. Springer, Berlin, pp 177–186
 65. Salge C, Ay N, Polani D, Prokopenko M (2015) Zipf's law: balancing signal usage cost and communication efficiency. *PLOS One* 10(10):e0139475
 66. Schumacher J, Wunderle T, Fries P, Jäkel F, Pipa G (2015) A statistical framework to infer delay and direction of information flow from measurements of complex systems. *Neural Comput* 27(8):1555–1608
 67. Schwarz G (1978) Estimating the dimension of a model. *Ann Stat* 6(2):461–464
 68. Shoham Y, Leyton-Brown K (2008) Multiagent systems: algorithmic, game-theoretic, and logical foundations. Cambridge University Press, Cambridge
 69. Sporns O, Chialvo DR, Kaiser M, Hilgetag CC (2004) Organization, development and function of complex brain networks. *Trends Cognit Sci* 8(9):418–425
 70. Stone P, Riley P, Veloso M (2000) The CMUnited-99 champion simulator team. In: Veloso M, Pagello E, Kitano H (eds) RoboCup-99: Robot Soccer World Cup III, vol 1856. Lecture notes in artificial intelligence. Springer, Berlin, pp 35–48
 71. Stone P, Riley P, Veloso M (2000) Defining and using ideal teammate and opponent models. In: Proc. of the 12th annual conf. on innovative applications of artificial intelligence
 72. Stone P, Veloso M (1998) Towards collaborative and adversarial learning: a case study in robotic soccer. *Int J Hum Comput Stud* 48(1):83–104
 73. Stone P, Veloso M (1999) Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artif Intell* 110(2):241–273
 74. Stone P, Veloso M (2000) Multiagent systems: a survey from a machine learning perspective. *Auton Robot* 8(3):345–383
 75. Taylor ME, Stone P (2009) Transfer learning for reinforcement learning domains: a survey. *J Mach Learn Res* 10(1):1633–1685
 76. Venables WN, Ripley BD (2003) Modern applied statistics with S, 4th edn. Springer, New York
 77. Vicente R, Wibrál M, Lindner M, Pipa G (2011) Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *J Comput Neurosci* 30(1):45–67
 78. Vilar L, Araujo D, Davids K, Bar-Yam Y (2013) Science of winning soccer: emergent pattern-forming dynamics in association football. *J Syst Sci Complex* 26:73–84
 79. Vrieze S (2012) Model selection and psychological theory: a discussion of the differences between the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). *Psychol Methods* 17(2):228–243
 80. Wasserman S, Faust K (1994) Social network analysis. Cambridge University Press, New York
 81. Weiss G (1999) Multiagent systems: a modern approach to distributed artificial intelligence. MIT press, Cambridge
 82. Whiteson S, Kohl N, Miikkulainen R, Stone P (2005) Evolving keepaway soccer players through task decomposition. *Mach Learn* 59(1):5–30
 83. Wu F, Zilberstein S, Chen X (2011) Online planning for multi-agent systems with bounded communication. *Artif Intell* 175:487–511
 84. Xu Z, Fitch R, Underwood J, Sukkarieh S (2013) Decentralized coordinated tracking with mixed discrete-continuous decisions. *J Field Robot* 30(5):717–740
 85. Zhang H, Chen X (2014) The decision-making framework of WrightEagle, the RoboCup 2013 soccer simulation 2D league champion team. In: Behnke S, Veloso M, Visser A, Xiong R (eds) RoboCup 2013: Robot Soccer World Cup XVII, vol 8371. Lecture notes in artificial intelligence. Springer, Berlin, pp 114–124