

Quantifying Long-Range Interactions and Coherent Structure in Multi-Agent Dynamics

Oliver M. Cliff^{*,**}

Joseph T. Lizier[†]
University of Sydney

X. Rosalind Wang[‡]

Peter Wang[‡]
CSIRO

Oliver Obst^{‡,§}

CSIRO
Western Sydney University

Mikhail Prokopenko[†]

University of Sydney

Abstract We develop and apply several novel methods quantifying dynamic multi-agent team interactions. These interactions are detected information-theoretically and captured in two ways: via (i) directed networks (interaction diagrams) representing significant coupled dynamics between pairs of agents, and (ii) state-space plots (coherence diagrams) showing coherent structures in Shannon information dynamics. This model-free analysis relates, on the one hand, the information transfer to responsiveness of the agents and the team, and, on the other hand, the information storage within the team to the team's rigidity and lack of tactical flexibility. The resultant interaction and coherence diagrams reveal implicit interactions, across teams, that may be spatially long-range. The analysis was verified with a statistically significant number of experiments (using simulated football games, produced during RoboCup 2D Simulation League matches), identifying the zones of the most intense competition, the extent and types of interactions, and the correlation between the strength of specific interactions and the results of the matches.

Keywords

Multi-agent dynamics, distributed computation, implicit communication, information storage, information transfer

1 Introduction

Multi-agent dynamics in complex biological or technological systems typically involves distributed communications and control, and its rigorous analysis becomes increasingly relevant to fundamental studies of artificial life and lifelike systems. In general, the challenges of distributed control are brought about by shared collective objectives of a multi-agent system (team), in which multiple autonomous agents must cooperate in making distributed decisions in optimizing the overall team objective [40]. In addition, not all communications occur explicitly within well-defined channels. Instead, complex multi-agent behaviors involve tacit interactions that can be characterized by implicit communications, including spatially long-range interactions with indirect effects. These implicit interactions need to be properly accounted for within specific feedback control loops.

* Contact author.

** Australian Centre for Field Robotics, The Rose Street Building J04, University of Sydney, NSW 2006, Australia. E-mail: o.cliff@acfr.usyd.edu.au

† Centre for Complex Systems, Civil Engineering Building J05, University of Sydney, NSW 2006, Australia. E-mail: joseph.lizier@sydney.edu.au (J.T.L.); mikhail.prokopenko@sydney.edu.au (M.P.)

‡ CSIRO, Locked Bag 17, North Ryde, NSW 1670, Australia. E-mail: rosaling.wang@data61.csiro.au (X.R.W.); peter.wang@data61.csiro.au (P.W.)

§ Centre for Research in Mathematics, School of Computing, Engineering and Mathematics, Western Sydney University, Locked Bag 1797, Penrith NSW 2751, Australia. E-mail: o.obst@westernsydney.edu.au

Furthermore, multi-agent dynamics are often constrained by either changing and partially unknown environmental factors or competing objectives of some adversaries engaged in directly opposing activities. This is typical, for example, in various team sports where some of the observable interactions cannot be simply reduced to algorithmic details of the agents, being affected by a multiplicity of concurrent activities.

Many team games, real and virtual, include rich interactions occurring dynamically and shaping the course of the contest both locally and globally. While the interactions within a team are usually constrained by cooperatively shared plans and tactical schemes, the interactions across the teams are created by opposing objectives of competing players. Generally, the interactions vary in strength over time and/or space, manifesting some tacit correlations that often are delayed in time and/or are long-range over the playing field.

Thus, distributed control of a multi-agent system, deployed in an adversarial environment, demands new techniques for identifying possibilities and features of feedback control loops. For instance, changing team tactics during a contest requires the team to quickly and coherently detect emergent patterns and regularities, quantify their strength and extent, and evaluate their potential impact on the overall performance.

While explicit communication between agents has been effectively analyzed in robotic information-gathering tasks [23], detection and quantification of implicit indirect interactions in distributed systems remains a challenge. This is primarily due to inaccessibility of the logic and neural processing of the opposing players as well as noise in the environment, highlighting the need for generic model-independent information-theoretic techniques.

Quantitative analysis of information-processing attributes of swarming in particular is a rapidly expanding cross-disciplinary field, ranging from biology [12] to statistical mechanics [6] to swarm engineering [30, 45]. For instance, information transfer in a swarm of fish was quantified as the normalized angular deviation of group direction, showing that transfer of information and decision making can occur in an animal group without explicit signals or individual recognition [12]. The maximum-entropy model was used to establish that local pairwise interactions between birds are sufficient to correctly predict the propagation of order throughout entire flocks of starlings [6]. An intuitive measure of information flow was used to identify behavioral strategy within simulated swarms, demonstrating swarm plasticity in response to changing environments [45].

For instance, information cascades within a simulated swarm were quantified by considering dynamic synchrony in collective motion of swarm individuals that do not exchange explicit messages [41, 68]. Specifically, the swarm's collective communications were captured using conditional transfer entropy [68]. A follow-up study compared such collective communications within two different swarms, one of which had a constraint imposed on the speed of its individuals [41]; and reported that the constrained swarm generated weaker information cascades and had more difficulties in self-organizing into a coherent state. We build on such advances here to detect and analyze implicit interactions within and between teams which are undertaking a specific collective task.

1.1 Networks for Multi-agent Dynamics

The first problem addressed in this study is identifying interaction networks that link together autonomous agents, without reconstructing the agents' logic and neural processing and using only the observational data, such as positional (e.g., planar) coordinates and their changes. The problem is difficult in that some of the dependences between agents are not discernible simply by correlating their corresponding locations over time—one needs to take into account the possibly directed nature of such correlations, where dynamics of one of the agents affects the positioning of another [10].

Quantitative analysis, in particular using complex systems theory, is increasingly being used in team sports to better understand and evaluate performance [1, 67] and identify networks between players. One of the recent examples is described by Fewell et al. [18], who analyzed basketball games as networks, where players are represented as nodes and passing density as edge weights: 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, and attempts to determine whether differences in team offensive strategy can be assessed by their network properties. Strategic networks considered by Fewell et al. include only explicit interactions (such as passes) within a team, and not implicit or spatially long-range interactions across teams.

Similar analysis was applied in the context of football as well, using passing data made available by Fédération Internationale de Football Association (FIFA) during the 2010 World Cup [44]. The study constructed a static weighted directed graph for each team (the *passing network*), with vertices corresponding to players and edges to passes, in order to provide a direct visual inspection of a team's strategy. The passing network was visualized by placing the nodes in positions roughly corresponding to the players' formation on the pitch, and enabling inspection of play patterns, hot spots, and potential weaknesses. Using different centrality measures, the study also determined the relative importance of each player in the game. This work—as well as the previous study of Duch et al. [14], which constructed and analyzed networks with one node for shots on target and one for wide shots—is limited to static passing networks, and again does not reveal spatially long-range interactions across teams.

Importantly, the multi-player dynamics of a football game was recently shown to exhibit self-similarities in the time evolution of player and ball positioning [25]. Specifically, the persistence time below which self-similarity holds has been estimated to be a few tens of seconds, implying that the volatility of football dynamics is an intrinsic feature of these games. Taking such volatility into account, the investigation by Vilar et al. [67] proposed a novel method of analysis that captures how teams occupy subareas of the field as the ball changes location. This study was important in focusing on the local dynamics of team collective behavior rather than individual player capabilities: When applied to football (soccer) matches, the method suggested that players' numerical dominance in some local subareas is a key to “defensive stability” and “offensive opportunity.” While the method rigorously used an information-theoretic approach (e.g., the uncertainty of the team numerical advantage across subareas was determined using the Shannon entropy), it was not aimed at and did not produce interaction networks, either explicit or implicit.

In our study we use *information dynamics*: a recent methodology for analysis of complex systems in general, including delayed and long-range effects in particular. It investigates the phenomenon of information processing (or computation) in a systematic way, by uncovering and quantifying information-theoretic roots of the most basic computational primitives: *storage*, *transmission*, and *modification of information* [34–37]. For instance, recent studies by Wang et al. [68] quantitatively verified the hypothesis that the collective memory within a swarm can be captured by *active information storage*. Higher values of storage are associated with higher levels of dynamic coordination. Furthermore, cascading information waves that correspond to long-range communications are captured by *conditional transfer entropy* [34, 35]. In other words, information transfer was shown to characterize the communication aspect of collective computation distributed within the swarm.

1.2 Coherent Structure in Distributed Communications

The second problem we address is classifying coherent dynamical situations within the multi-agent games, in the context of distributed communications. For example, during a game, each player (dependent on his tactical role) is engaged in dynamics that are affected both by (i) the player's history of actions (persistence or rigidity), and (ii) spatially long-range effects of other players' actions (sensitivity or responsiveness). Therefore, we may want to form an abstract state space with variables quantifying these features, and consider a structure of this space, aiming to classify the games and game situations by identifying coherent regions within the space.

In general, one of the defining features of complex computation is a coherent information structure understood as some pattern or configuration appearing in a state space formed by information-theoretic quantities, such as transfer entropy [56] and excess entropy [13]. The information dynamics state space diagrams are known to provide insights that are not immediately visible

when the measures are considered in isolation [36]. One example is a structure for a class of systems (such as logistic maps) that can be examined by plotting *average* excess entropy versus entropy rate while changing a system parameter [17]. Another example is a characterization of complexity of distributed computation within the spatiotemporal dynamics of cellular automata (CA) via state space diagrams formed by transfer entropy and active information of the CA rules [36]. In this example, each point in the state space quantifies both the communication and the memory operations of a cellular automaton.

Consequently, in addition to identifying implicit interaction networks, we intend to adopt the methods of coherent information structure in classifying repeatable collective dynamics in game situations. In doing so, we shall use game dynamics produced within the RoboCup environment.

1.3 The RoboCup Initiative

During the last two decades, the RoboCup initiative has essentially superseded chess as a benchmark for artificial intelligence (AI). 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 AI, robotics, computer vision, and several other related areas [27], with the overarching RoboCup goal of developing a team of humanoid robots capable of defeating the FIFA World Cup champion team (the “Millennium Challenge”). From the outset of the RoboCup effort it was recognized that RoboCup is different from the previous benchmark (chess) in several crucial elements: environment (static versus dynamic), state change (turn-taking versus real-time), information accessibility (complete versus incomplete), sensor readings (symbolic versus non-symbolic), and control (central versus distributed) [4]. Since 1997, this ambitious goal has been pursued along two general complementary paths [26]: the physical robot league, and the software agent (simulation) league [43].

The RoboCup 2D Soccer Simulation League specifically targets the research question of how the optimal collective dynamics can result from autonomous decision making under constraints set by tactical plans and teamwork (collaboration) as well as opponents (competition) [8, 28, 42, 50–52, 54, 59, 64, 69]. In answering this question it becomes important to measure the mechanisms for, and to discover the patterns of, dynamic spatiotemporal interactions between different players. In this article we describe our approach to detection and quantification of dynamic interactions in simulated football games, produced during RoboCup 2D Simulation League matches.

1.4 Contributions

The contributions of this article are threefold. Firstly, we further validate our information-theoretic analysis of the RoboCup-2012 2D simulation [10] with the most recent data sets from RoboCup-2014 and new relative measures. Secondly, we extend the analysis towards data sets that reveal *asymmetric interaction networks*, which, for the first time, discern and quantify a number of asymmetries in the tactical schemes used by the teams. Thirdly, we produce novel state space plots (coherence diagrams), which classify some repeatable multi-agent situations within the games, by contrasting persistence of players’ actions with sensitivity to behavior of other players.

2 Technical Preliminaries

In this section, we introduce two fundamental information-theoretic quantities: *active information storage* and *transfer entropy*. These measures are considered for stochastic temporal processes X , that is, sequences of random variables $(\dots, X_{n-1}, X_n, X_{n+1}, \dots)$ with associated realizations $(\dots, x_{n-1}, x_n, x_{n+1}, \dots)$ for countable time indices n .

The active information storage \mathcal{A}_X quantifies the information storage component that is directly in use in the computation of the next value of a process [37]. More precisely, the active information

storage is the average mutual information between the (semi-infinite) past state of the process and its next value:

$$\mathcal{A}_X = \lim_{\ell \rightarrow \infty} \left\langle \log_2 \frac{p(x_n^{(\ell)}, x_{n+1})}{p(x_n^{(\ell)})p(x_{n+1})} \right\rangle, \tag{1}$$

where we have introduced the shorthand notation $x_n^{(\ell)} = \{x_{n-\ell+1}, \dots, x_{n-1}, x_n\}$. In practice, one deals with finite- ℓ estimates $\mathcal{A}_X(\ell)$.

The transfer entropy is designed to detect asymmetry in the interaction of subsystems by distinguishing between *driving* and *responding* elements [56]. Specifically, transfer entropy captures information transmission from a source Y to a destination X process as the average information provided by the source variable Y_n about the next destination variable X_{n+1} in the context of the past state of the destination $X_n^{(\ell)}$ [34, 56].¹ The transfer entropy is computed as

$$T_{Y \rightarrow X} = \lim_{\ell \rightarrow \infty} \left\langle \log_2 \frac{p(x_{n+1} \mid x_n^{(\ell)}, y_n)}{p(x_{n+1} \mid x_n^{(\ell)})} \right\rangle. \tag{2}$$

Again, in practice one deals with finite- ℓ estimates $T_{Y \rightarrow X}(\ell)$.

It is important to realize that information transfer between two variables does not require an explicit communication channel; it rather indicates a high degree of directional synchrony or nonlinear correlation between the source and the destination. It characterizes a degree of *predictive* information transfer, that is, “if the state of the source is known, how much does that help to predict the state of the destination?” [34].

Sometimes it is useful to condition the local information transfer on another contributing process Z , considering the *conditional transfer entropy* [34, 35]:

$$T_{Y \rightarrow X|Z} = \lim_{\ell \rightarrow \infty} \left\langle \log_2 \frac{p(x_{n+1} \mid x_n^{(\ell)}, y_n, z_n)}{p(x_{n+1} \mid x_n^{(\ell)}, z_n)} \right\rangle, \tag{3}$$

and in practice we deal with finite- ℓ estimates $T_{Y \rightarrow X|Z}(\ell)$. We may also utilize temporally local values in order to trace the information dynamics over time—for example, identifying its peaks during specific moments (see [32]). In this article we employ the measures with a constant embedding dimension, and henceforth omit the argument ℓ unless necessary.

3 Tactical Information Dynamics

In order to estimate the strength of directed coupling between two agents, we compute the average transfer entropy between them during any given game. In this section we formally define the state space of a game and describe the information-theoretic measures above in this context as classifying the responsiveness and rigidity of a team.

We shall be using the notion of a tactical formation that describes how the players in a team are generally positioned on the field in terms of their roles (the number of: defenders-midfielders-attackers), for example, the 4-3-3 formation with four defenders, three midfielders, and three forwards. Of course, during a game, any player may be drawn to a position fairly remote from its area of responsibility defined by the role (for instance, a defender may join a particular attack), but in general the players tend to stay within their distinct areas, specified by some prior configurations and/or distinguishable by spatial pattern matching; see Figure 1. Hence, any dynamic coherence observed

¹ In general, one can consider a sequence of source variables $Y_n^{(l)}$; however, we only consider single variables, because Y_n is assumed directly causal to X_{n+1} . See further discussion in [34].




Figure 1. Motion trace diagram. A trace curve represents the motion of the left midfielder (player 7) of the left team, during an entire game (solid yellow for regular play-on time points, and dotted black for non-play-on times, e.g., free kicks). The role of player 7 is distinguishable as the left midfielder.

in the motion of players that are spatially separated on the field due to their tactical roles (e.g., a midfielder of one team and an opponent’s defender) can be interpreted as spatially long-range implicit interactions.

A game g contains N time steps and is played between two teams \mathbb{X} and \mathbb{Y} with M agents each. The dynamics of the game is captured by the realization of two sets of stochastic processes $\mathcal{X} = \{X^1, \dots, X^M\}$ and $\mathcal{Y} = \{Y^1, \dots, Y^M\}$, that is, the movements of players in teams \mathbb{X} and \mathbb{Y} , respectively. The measurement of each temporal process X is therefore a sequence of positional data (x_1, \dots, x_N) ; in this article we consider observations x_n as the *change* in the 2D positional vector of the agent. Note that in this work we use the terms process, agent, and player interchangeably, depending on context.

3.1 Transfer Entropy as Player Responsiveness

For each game g , the transfer entropy is calculated between each source agent Y^i and destination agent X^j , in the context of some other dynamics Z , denoted $T_{Y^i \rightarrow X^j | Z}^g$. In the remainder of this article, the relative position of the ball is always conditioned upon in order to compute the transfer entropy in the context of the game, since this context is greatly affected by the ball trajectories in football matches. We also define the average transfer entropy over a range of source-destination pairs, targeting subsets $\mathcal{Y}_\alpha \subseteq \mathcal{Y}$ and $\mathcal{X}_\beta \subseteq \mathcal{X}$:²

$$T_{\mathcal{Y}_\alpha \rightarrow \mathcal{X}_\beta | Z}^g = \frac{1}{|\mathcal{Y}_\alpha| |\mathcal{X}_\beta|} \sum_{Y^i \in \mathcal{Y}_\alpha} \sum_{X^j \in \mathcal{X}_\beta} T_{Y^i \rightarrow X^j | Z}^g. \tag{4}$$

² We note a subtle distinction here: $T_{\mathcal{Y}_\alpha \rightarrow \mathcal{X}_\beta | Z}^g$ is not equal to the multivariate transfer entropy [33] from the set \mathcal{Y}_α to \mathcal{X}_β (conditioned on Z) as a whole in general, because of dependences within and across the sets. Were one wishing to measure such a multivariate transfer entropy, then $T_{\mathcal{Y}_\alpha \rightarrow \mathcal{X}_\beta | Z}^g$ could be viewed as an approximation to it (ignoring these dependences) in order to avoid dimensionality issues.

The average transfer entropy defined for specific subsets of team processes is useful in considering distributed communications across agents with specific roles (e.g., attackers and defenders in football).

Building upon the information dynamics measures, it is possible to investigate role-based behavior with complex interactions. In applying information dynamics to the RoboCup 2D Simulation League we use the following definition:

Definition 1. *The responsiveness of player X to player Y during the game g is defined as the information transfer $T_{Y \rightarrow X|Z}^g$ from the source Y (e.g., dynamics of player Y) to the destination X (e.g., dynamics of another player X) in the context of some other dynamics Z (e.g., the movement of the ball).*

That is, the destination player X responds—for example, by repositioning—to the movement of the source player Y . This may apply to many situations on the field. For instance, when one team’s forwards are trying to better avoid their opponent’s defenders, we consider the information transfer $T_{Y^i \rightarrow X^j|Z}^g$ from defender-agent processes $Y^i \in \mathcal{Y}_d$ to forward-agent processes $X^j \in \mathcal{X}_a$, where the roles of the agents are determined by their placements in a given tactical formation. Henceforth, we omit the game index g and the condition variable B when there is no ambiguity. Vice versa, the dynamics of the opponent’s defenders, who are trying to better mark our team’s forwards, are represented in the information transfer $T_{X^j \rightarrow Y^i}$ from forward-agent processes $X^j \in \mathcal{X}_a$ to defender-agent processes $Y^i \in \mathcal{Y}_d$. These two examples specifically consider a coupling between the attack line \mathcal{X}_a of our team and the defense line \mathcal{Y}_d of the opponent’s team.

3.2 Active Information Storage as Player Rigidity

Our analysis also involves computation of the active information storage within the teams. We can define the average active information storage over a range of agents in a game g , targeting subsets $\mathcal{X}_\beta \subseteq \mathcal{X}$:³

$$A_{\mathcal{X}_\beta}^g = \frac{1}{|\mathcal{X}_\beta|} \sum_{X^i \in \mathcal{X}_\beta} A_{X^i}^g. \tag{5}$$

We characterize a team’s rigidity $\mathcal{A}_\mathcal{X}$ as the average of information storage values for all players of the team, according to the following definition.

Definition 2. *The rigidity of a player’s dynamics is defined as the information storage A_X within a process X (e.g., the dynamics of player X).*

The average information storage, or rigidity, within a team $\mathcal{A}_\mathcal{X}$ is high whenever one can predict the motion of some players based on the movements of their past. In these cases, the players are not as independent of their previous movements as a complex or swarm behavior may warrant, making the dynamics less versatile.

How much does a team’s rigidity and responsiveness contribute to a game’s scoreline? To answer this question, one can analyze the correlation between a number of measures and the scoreline $\delta S^g = S_\mathcal{X}^g - S_\mathcal{Y}^g$, where S^g is the number of goals scored by each team.

The utilized measures are relative; for example, the relative team responsiveness δT is calculated by comparing the transfer from team \mathcal{Y} to team \mathcal{X} with the transfer in the other direction, that is, $\delta T^g = T_{\mathcal{Y} \rightarrow \mathcal{X}|Z}^g - T_{\mathcal{X} \rightarrow \mathcal{Y}|Z}^g$. Table 1 summarizes different relative measures, specified for different tactical roles in a typical football formation. We should like to point out that we introduce here

³ As per footnote 2, $A_{\mathcal{X}_\beta}^g$ is not equal to the collective active information storage as defined for the multivariate set \mathcal{X}_β in general, due to dependences between the variables. Again, were one wishing to measure such a collective quantity, then $A_{\mathcal{X}_\beta}^g$ could be seen as an approximation to it (ignoring these dependences), which avoids dimensionality issues.

Table I. Tactical information dynamics measures.

Primitive	Metric	Equation	Description
Transmission	$\delta T(\mathcal{X}, \mathcal{Y})$	$T_{\mathcal{Y} \rightarrow \mathcal{X}} - T_{\mathcal{X} \rightarrow \mathcal{Y}}$	Team relative responsiveness (RR)
	$\delta T_{a \rightarrow d}(\mathcal{X}, \mathcal{Y})$	$T_{\mathcal{Y}_a \rightarrow \mathcal{X}_d} - T_{\mathcal{X}_a \rightarrow \mathcal{Y}_d}$	RR of defenders to opponent attackers
	$\delta T_{m \rightarrow m}(\mathcal{X}, \mathcal{Y})$	$T_{\mathcal{Y}_m \rightarrow \mathcal{X}_m} - T_{\mathcal{X}_m \rightarrow \mathcal{Y}_m}$	RR of midfielders to opponent midfielders
	$\delta T_{d \rightarrow a}(\mathcal{X}, \mathcal{Y})$	$T_{\mathcal{Y}_d \rightarrow \mathcal{X}_a} - T_{\mathcal{X}_d \rightarrow \mathcal{Y}_a}$	RR of attackers to opponent defenders
Storage	$\delta A(\mathcal{X}, \mathcal{Y})$	$A_{\mathcal{X}} - A_{\mathcal{Y}}$	Relative team rigidity
	$\delta A_d(\mathcal{X}, \mathcal{Y})$	$A_{\mathcal{X}_d} - A_{\mathcal{Y}_d}$	Relative defender rigidity
	$\delta A_m(\mathcal{X}, \mathcal{Y})$	$A_{\mathcal{X}_m} - A_{\mathcal{Y}_m}$	Relative midfielder rigidity
	$\delta A_a(\mathcal{X}, \mathcal{Y})$	$A_{\mathcal{X}_a} - A_{\mathcal{Y}_a}$	Relative attacker rigidity

new relative measures, expanding on the ones analyzed in [10]. Specifically, the previous study [10] compared attacking versus defending lines, that is, analyzed $T_{\mathcal{Y}_a \rightarrow \mathcal{X}_d} - T_{\mathcal{X}_d \rightarrow \mathcal{Y}_a}$, while in this work we compare attacking versus attacking lines $T_{\mathcal{Y}_d \rightarrow \mathcal{X}_a} - T_{\mathcal{X}_d \rightarrow \mathcal{Y}_a}$ on the one hand, and defending versus defending lines $T_{\mathcal{Y}_a \rightarrow \mathcal{X}_d} - T_{\mathcal{X}_a \rightarrow \mathcal{Y}_d}$ on the other hand. This change is addressing a different question of evaluating relative performance of a specific tactical line (role). Note also that we use averaged pairwise calculations in (4) and (5), as opposed to a multivariate approach (as in [33, 38]). These two approaches are only equivalent if the individual player processes are independent.

4 Interaction Diagrams

We describe here another information dynamics tool, interaction diagrams, which provide a simplified view of the strongest pairwise interactions into (Section 4.1) or out from (Section 4.2) each agent.

4.1 Information-Sink Diagrams

Once the game's average transfer entropy $T_{Y^i \rightarrow X^j | Z}^g$ is determined for each pair (Y^i, X^j) , we identify the *source* opposing agent $\hat{y}_{\mathcal{Y}}(j)$ described by the process $Y^i \in \mathcal{Y}$ that transfers maximal information to process X^j for the given agent j :

$$\hat{y}_{\mathcal{Y}}(j, g) = \arg \max_{1 \leq k \leq M} T_{Y^k \rightarrow X^j | Z}^g. \quad (6)$$

Over a number of games G , we select the source agent $\hat{y}_{\mathcal{Y}}(j)$ that transfers maximal information to X^j most frequently, as the mode of the set $\{\hat{y}_{\mathcal{Y}}(j, 1), \dots, \hat{y}_{\mathcal{Y}}(j, G)\}$. Then, we consider the average information transfer between these two processes $Y^i = Y^{\hat{y}_{\mathcal{Y}}(j)}$ and X^j across all games:

$$T_{Y^i \rightarrow X^j | Z} = \frac{1}{G} \sum_{g=1}^G T_{Y^i \rightarrow X^j | Z}^g. \quad (7)$$

Intuitively, the movement of the source agent $\hat{y}_{\mathcal{Y}}(j)$ affected the agent j more than the movement of any other agent in team \mathbb{Y} . That is, the agent j was responsive most to movement of the source

agent $\hat{y}(j)$. Crucially, when we use the notion of *responsiveness* to another (source) agent, we do not load it with such semantics as being dominated by, or driven by, that other agent. Higher responsiveness may in fact reflect either useful reaction to the opponent’s movements (e.g., good marking of the source), or a helpless behavior (e.g., constant chase after the source). Vice versa, generating a high responsiveness from another agent may result in either a useful dynamic (e.g., positional or even tactical dominance over the responding agent), or a wasteful motion (e.g., being successfully marked by the responding agent). In short, responsiveness captured in the maximal transfer $T_{Y^i \rightarrow X^j|Z}$ detects a directed coupling from the source process Y^i to the responding process X^j and at face value alone should not be interpreted in general as a simple index for comparative performance. It is, however, a useful identifier of the opponents’ source player that was affecting a given agent j most.

Given a series of games, we identify the source-responder pairs by finding the source agent for each of the agents on both teams (always choosing the source among the opponents). The pairs $(\hat{y}(j), j)$ identified for each agent j in team \mathbb{X} treated as a destination are combined in an *information-sink diagram* $\hat{D}(\mathcal{Y}, \mathcal{X}) = \{(\hat{y}(j\mathcal{X}), j\mathcal{X}) \mid \forall j\mathcal{X} \in [1, M]\}$. The interaction diagram $\hat{D}(\mathcal{Y}, \mathcal{X})$ visualizes a directed graph with $2m$ nodes representing players, and with the edges representing all source-responder pairs, where a single edge is incoming to every agent from the corresponding source. One may extend the diagrams by specifying the weight of each edge with the corresponding transfer entropy.

4.2 Information-Source Diagrams

Similarly, having obtained the average transfer entropy during a game, $T_{Y^i \rightarrow X^j|Z}^g$, for all pairs, we identify the *responder* agent $\check{y}_\mathcal{X}$ described by the process $X^j \in \mathcal{X}$ that received maximal information from process Y^i for the given agent i . Formally, for any game g ,

$$\check{y}_\mathcal{X}(i, g) = \arg \max_{1 \leq k \leq M} T_{Y^i \rightarrow X^k|Z}^g. \tag{8}$$

Over a number of games G , we select the responder agent $\check{y}_\mathcal{X}(i)$ to which maximal information was transferred by Y^i most frequently, as the mode of the series $\{\check{y}_\mathcal{X}(i, 1), \dots, \check{y}_\mathcal{X}(i, G)\}$. Finally, we consider the average information transfer between these two processes Y^i and $X^j = X^{\check{y}_\mathcal{X}(i)}$ across all games:

$$T_{Y^i \rightarrow X^j|Z} = \frac{1}{G} \sum_{g=1}^G T_{Y^i \rightarrow X^j|Z}^g. \tag{9}$$

The pairs $(i, \check{y}_\mathcal{X}(i))$ identified for each agent i in team \mathbb{Y} treated as a source are combined in an *information-source diagram* $\hat{D}(\mathcal{Y}, \mathcal{X}) = \{(i\mathcal{Y}, \check{y}_\mathcal{X}(i\mathcal{Y})) \mid \forall i\mathcal{Y} \in [1, M]\}$.

The intuition in this case is the same as in the previous subsection—the difference is that now we identify the highest responder agent, having selected a source. In general, the agent i in team \mathbb{Y} may be the most informative source for the agent j in team \mathbb{X} , but the agent j may be not the best responder to the agent i among all possible responders in team \mathbb{X} , and vice versa.

While an information-sink diagram reflects more where the information tends to be transferred to, an information-source diagram tends to depict where the information is transferred from.

4.3 Information-Sink and -Source Diagrams as Efficient Simplifications

Neither of the diagrams presents a complete story, each highlighting only a small part of the overall information dynamics. There are more comprehensive network diagrams, specifically known as *effective networks* or *effective connectivity networks*, which seek to infer a circuit model that can replicate and indeed explain the time series of the nodes in the network [20, 57]. Such effective network

inference is popular in analysis of data sets obtained from neural recordings, and the transfer entropy is a well-utilized tool in this area (e.g., see [11, 15, 16, 38, 39, 58, 61, 66, 70]). Effective networks may be constructed for example with the edges representing in descending order the highest information transfers for all the pairs, retaining a given number of such links, keeping the edges for the information amounts above a certain threshold, considering higher-order interactions to remove redundant links and include synergistic effects, and so on—in these instances, some agents may have no incoming or outgoing links at all. Furthermore, we note that information-sink and -source diagrams ignore interactions within teams, and of course both these and full effective networks represent observational correlations rather than strict causation (by specifically using a Wiener-Granger interpretation of causality).

Nevertheless, we believe that the interaction diagrams presented here are valuable, as a simplified view of the full effective network representation of the set of agents influencing and influenced by each other agent: They are particularly simple and easy to interpret, and crucially are computationally efficient. Specifically, for an information-sink diagram every agent has an incoming edge, and for an information-source diagram every agent has an outgoing edge, representing the strongest respective in- or outgoing interactions for that agent. Also, these diagrams provide a significantly more efficient analysis than full effective network inference, computing only $O(M^2)$ transfer entropies rather than additionally examining higher-order interactions, and avoiding additional computations for statistical significance measurements. Such efficiency is a particularly important consideration if such a method is to be used online during RoboCup games in the future.

5 State Space Coherence Diagrams

The study by Lizier et al. [36] diagrammatically demonstrated that more coherent structures in state space plots can be observed in systems (cellular automata) with higher degrees of complexity. Motivated by methods described in [36], we investigate coherent information structures observed as patterns in a state space formed by tactical information dynamics measures, aiming to reveal structure in the relationship between the team’s rigidity and responsiveness. The positional dynamics of each agent depends in general on its tactical role in the game and is quantified by its responsiveness (measured by information transfer) and rigidity (measured by information storage). These two measures will specifically be used in forming the two-dimensional state space, where relative responsiveness is plotted as a function of relative rigidity (see Table 1 for definitions).

Identifying a coherent structure in the relationship between responsiveness and rigidity allows us to classify coherent dynamic situations in the context of distributed communications. For example, the dynamics of agents performing in a specific role, such as attackers, may be characterized by both lower rigidity and lower responsiveness to the opponent’s defenders than in the dynamics of other agents. Coherence diagrams are intended to visualize such dynamic clustering in the state space formed by the corresponding information-theoretic measures. Furthermore, once these dynamic clusters are highlighted as subregions of the space, it is possible to “zoom in” by considering correlations of the points within these regions with the scorelines of the corresponding games, and identifying which regions (clusters) map to more successful games.

In particular, we introduce two different state space plots (coherence diagrams) intended to capture different spatiotemporal interactions across teams: (1) tactical information dynamics in relation to tactical roles (defender, midfielder, attacker), and (2) information dynamics partitions correlated with the scorelines. The state space diagrams for each team are produced by computing the following state space points: $(\delta A_{ab}, \delta T_{a \rightarrow a})$, $(\delta A_{mb}, \delta T_{m \rightarrow m})$, and $(\delta A_{ab}, \delta T_{d \rightarrow a})$. Then the first coherence diagram is obtained by plotting these points on the respective axes with a distinct color for each tactical role, viz., defenders, midfielders, and attackers (see Figure 6 in Section 6.3). Another coherence diagram is given for each tactical role by selecting the points corresponding to the roles, and colour-mapping them with the corresponding scoreline δS (see Figures 7 and 8 in Section 6.3).

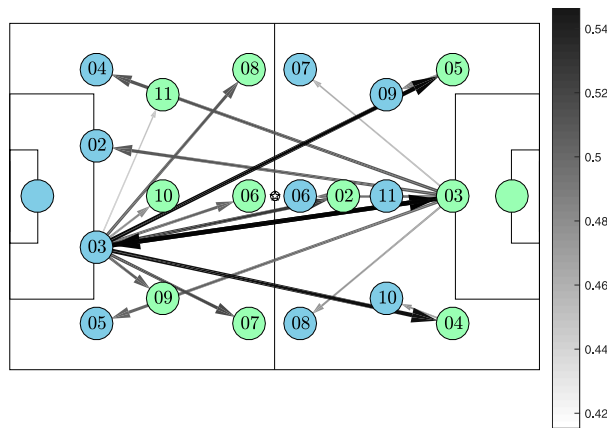
6 Results and Discussion

To compute the measures described in previous sections and produce interaction diagrams and state space coherence diagrams, we carried out multiple iterative experiments matching team Gliders [53] up against teams Cyrus [24] and HELIOS [3], denoted by \mathcal{G} , \mathcal{C} , and \mathcal{H} , respectively. Gliders was the runner-up (vice-champion) team for RoboCup-2014, while HELIOS and Cyrus were fourth- and fifth-ranked teams.

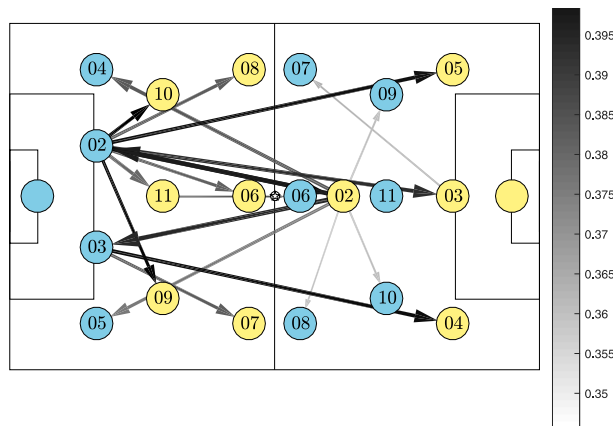
All information-theoretic measures were computed using the JIDT toolkit [31], with finite history lengths $k = 1$. For the information-sink and base diagrams, kernel estimation was used with a kernel width of 0.4 standard deviations of the data. For the state space coherence diagrams, Kraskov-Stögbauer-Grassberger estimation [22, 29] was used with $K = 4$ nearest neighbors.

6.1 Interaction Diagrams

Figure 2 presents the information-sink interaction diagrams, built over 400 games between Gliders and Cyrus $\hat{D}(\mathcal{C}, \mathcal{G})$ and Gliders and HELIOS $\hat{D}(\mathcal{H}, \mathcal{G})$. Analogously, Figure 3 shows the information-source

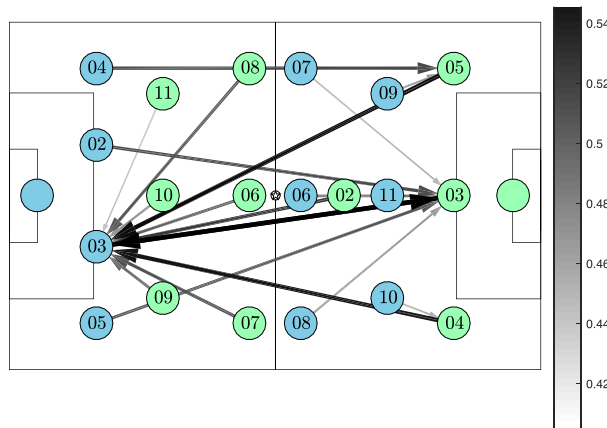


(a) Information-sink diagram for Gliders (blue) and Cyrus (green).

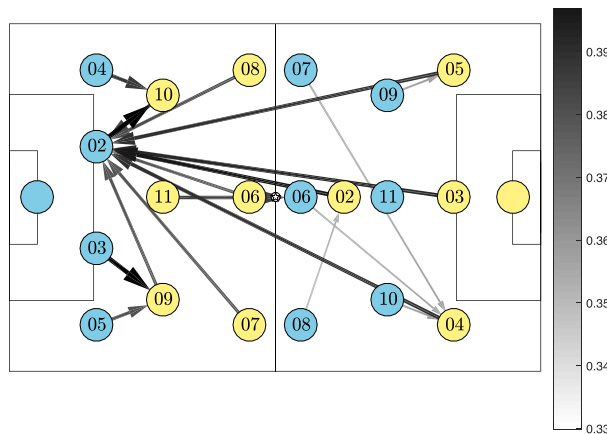


(b) Information-sink diagram for Gliders (blue) and HELIOS (yellow).

Figure 2. Interaction-sink diagrams. Arrows represent highest information transfer between players. Grayscale colormap is used to indicate the strength of transfer, varying smoothly from white (weakest) to black (strongest). Example of the most pronounced interactions: All Cyrus players strongly respond to the motion of the right center-back of Gliders (player 03), indicating the strong asymmetry of Cyrus tactics in preferring to play on the left wing.



(a) Information-source diagram for Gliders (blue) and Cyrus (green).



(b) Information-source diagram for Gliders (blue) and HELIOS (yellow).

Figure 3. Interaction-source diagrams. Arrows represent highest information transfer between players. Grayscale is used to indicate the strength of transfer, varying smoothly from white (weakest) to black (strongest). Example of the most pronounced interactions: All HELIOS players strongly drive the left center-back of Gliders (player 02), indicating the strong asymmetry of HELIOS dynamics in preferring to play on the right wing.

interaction diagrams built over the same 400 games between Gliders and Cyrus $\check{D}(\mathcal{C}, \mathcal{G})$ and Gliders and HELIOS $\check{D}(\mathcal{H}, \mathcal{G})$. The nodes in each diagram are shown in positions roughly corresponding to the players' formation on the field; for example, Gliders follow the 4-3-3 formation with four defenders playing line defense, three midfielders, and three attackers, whereas Cyrus and HELIOS utilize one of the defenders (the player with the number 02) as a defensive midfielder, thus loosely following 3-4-3 formation with four midfielders.

Perhaps most significantly, the interaction diagrams capture spatially long-range information transfer between agents. The transfer entropy accounts for information embedded in the source agent's actions by conditioning on prior states, thus allowing us to reconstruct diagrams that represent the implicit communication between source-destination pairs in a multi-agent system. In this way, we use transfer entropy to filter the observations of the agents and reveal latent structures in the multi-agent dynamics.

Several interesting observations can be made. To some extent, the interaction diagrams exhibit lateral symmetry, which is expected given the symmetric formations of the teams. However, and




Figure 4. A white curve traces the ball motion during an entire game between Gliders (left) and Cyrus (right). Note the asymmetry in Cyrus attacks, as well as a significant ball trace outside of the midfield.

perhaps more importantly, there are some clearly asymmetric connections. For example, the most pronounced interactions are observed with all Cyrus players strongly responding to the motion of the right center-back of Gliders (player 03), which reveals the strong asymmetry of Cyrus dynamics in preferring to play on their left wing. This is a feature that has been successfully exploited by Gliders in allocating suitable defensive resources on this wing, resulting in a statistically significant performance gain (an increase in the average goal difference from 1.55 ± 0.03 to 1.80 ± 0.02 , over more than 6000 games: an improvement of 16%). Similarly, all HELIOS players strongly “drive” the left center-back of Gliders (player 02), also highlighting the strong asymmetry of HELIOS dynamics in preferring to play on their right wing. Again, this can be tactically exploited.

In Figure 2a it is evident that the defenders are the most responsive of both teams, showing that the games between Gliders and Cyrus unfold outside the midfield; see Figure 4. On the other hand, Figure 2b reveals a more disordered responsiveness between the teams, indicating that a lot of interactions occur in midfield during the games between Gliders and HELIOS. We also point out that the highest information transfer value computed over the games between Gliders and HELIOS (≈ 0.4 bits in Figure 2b and Figure 3b) is less than the lowest value computed over the games between Gliders and Cyrus (≈ 0.42 bits in Figure 2a and Figure 3a). This means that the Gliders and HELIOS players are more independent in their respective motions on average.

Specifically, Gliders attackers mostly respond to Cyrus defenders, and Gliders midfielders and defenders mostly respond to the Cyrus central defender (player 03), which is typically moving across wider areas often playing a *sweeper* role.⁴

This coupling is similar to patterns observed in Gliders and HELIOS dynamics; however, the interactions are generally weaker and are spread amongst more players than just one central defender, because both HELIOS central defenders take an active part in defending the area.

⁴ The sweeper (or libero) is a more versatile center-back who “sweeps up” the ball if an opponent manages to breach the defensive line. This position is rather more fluid than that of other defenders who man-mark their designated opponents.

In summary, the findings demonstrate applicability of the information dynamics measures to analysis of the dynamics of multi-agent teams, revealing the player pairs with most intense interactions and the extent of the resultant dependences.

6.2 Correlation with Performance

In this subsection, we correlate measures of relative responsiveness (either tactical role-by-role or team overall), as well as rigidity, with the game scorelines, and identify the tactical roles that affected the games more. That is, we compute a correlation coefficient between a series of game scorelines and a series of information dynamics values for a game. For clarity, we discuss mainly the interpretation of the correlations in the context of the Gliders’ performance.

Table 2 presents the correlation coefficients between scorelines and various information-based measures that were summarized in Table 1. Generally, the observed correlations are consistent for all measures across both opponent teams, with the exception of $\delta T_{d \rightarrow a}(\mathcal{G}, \mathcal{Y})$, which differ in sign. All of the correlations displayed in Table 2 are statistically significant at $p = 0.01$ (one-tailed test with Bonferroni correction for 16 comparisons). We begin our analysis with the measures based on information transfer.

Overall, the higher responsiveness of a team to all opponent players is detrimental to its winning chances, that is, the more responsive Gliders is on average to opponents, indicated by higher $\delta T(\mathcal{G}, \mathcal{Y})$, the less positive is the scoreline δS . However, when looking at tactical lines role by role (e.g., comparing the relative responsiveness of defenders to attackers between two teams), we observe in general the opposite effect: higher responsiveness is an indication of winning. In particular, if the Gliders defenders are more responsive to their immediate opposing line of Cyrus (or HELIOS) attackers than Cyrus (or HELIOS) defenders are to Gliders attackers, that is, $T_{y_d \rightarrow g_d} > T_{g_d \rightarrow y_d}$, then team Gliders has a higher chance of winning. Similarly, Gliders tends to win if its midfielders are more responsive than the midfield opposition (either Cyrus or HELIOS), that is, positive $\delta T_{m \rightarrow m}(\mathcal{G}, \mathcal{Y})$ can be used as a precursor for a winning prediction. This means that high relative responsiveness across tactical lines is indicative of a behavior positively contributing to the performance (e.g., defenders are successfully marking the opponent attackers, or midfielders are successfully finding open zones amongst opponent midfielders in anticipation of teammate passes using Voronoi diagrams [53]), while the overall high relative responsiveness across all players $\delta T(\mathcal{G}, \mathcal{Y})$ may suggest an adverse outcome, due to an excessive unstructured dependence on the opposition.

The relative responsiveness $\delta T_{d \rightarrow a}(\mathcal{G}, \mathcal{C})$ is negatively correlated with the scoreline, and deserves a separate explanation. The lower responsiveness $T_{c_d \rightarrow g_a} < T_{g_d \rightarrow c_a}$ means that Gliders attackers are less predictable in their response to the opponent defenders than Cyrus attackers, and this opens up more scoring opportunities. In other words, unpredictability of attackers’ motion is a positive feature characteristic of opportunity-seeking behavior, unlike the responsive tracking behavior of defenders, which are typically engaged in trying to actively mark the opponents attackers.

Table 2. Correlation of information dynamics measures with scoreline.

Team (\mathcal{Y})	$\delta T(\mathcal{G}, \mathcal{Y})$	$\delta T_{a \rightarrow d}(\mathcal{G}, \mathcal{Y})$	$\delta T_{m \rightarrow m}(\mathcal{G}, \mathcal{Y})$	$\delta T_{d \rightarrow a}(\mathcal{G}, \mathcal{Y})$
Cyrus (\mathcal{C})	-0.601	0.607	0.338	-0.427
HELIOS (\mathcal{H})	-0.466	0.455	0.616	0.211
	$\delta A(\mathcal{G}, \mathcal{Y})$	$\delta A_d(\mathcal{G}, \mathcal{Y})$	$\delta A_m(\mathcal{G}, \mathcal{Y})$	$\delta A_a(\mathcal{G}, \mathcal{Y})$
Cyrus (\mathcal{C})	-0.613	0.223	-0.558	-0.616
HELIOS (\mathcal{H})	-0.683	0.380	-0.703	-0.642




Figure 5. Motion trace diagram. A yellow-black trace curve represents the motion of the center-forward (player 11) of the left team (HELIOS) during an entire game (solid yellow for regular play-on time points, and dotted black for non-play-on times, e.g., free kicks). A white curve traces the ball motion during the game. Note that the majority of both traces lies in midfield.

The relative responsiveness for attackers $\delta T_{d-a}(\mathcal{G}, \mathcal{H})$ is still positively correlated with the scoreline, and should be interpreted in the context of the interaction diagrams, which indicated that in the games between Gliders and HELIOS most of the action occurred in midfield anyway, and so the attackers are mostly engaged in midfielder-like behavior, as can be seen in Figure 5. Hence it may be expected that the high relative responsiveness of attackers in this contest is still positively related to performance.

The rigidity of a team *as a whole* is also detrimental to its goal-scoring capabilities, as shown by the fact that δA is negatively correlated with the scorelines in both contests. This is an expected result for the team players, which are highly predictable with respect to their histories. An analysis for each role reveals that rigidity of either midfielders (δA_m) or attackers (δA_a) is also negatively correlated with performance. However, rigidity of defenders' movements (δA_d) is a positive feature across both matchups. This can be explained by a specific tactical behavior employed by the team Gliders: *line defense*, which is highly dependent on an ability to create offside traps by a simultaneous motion of all four defenders. This defensive tactic produces more synchronous actions and results in successful but predictable behavior of each player (on average), captured in turn by the players' rigidity. As long as this rigidity is not exploited by the opponents, the performance is likely to remain positively correlated.


The notion that the scoreline is correlated with a team's information dynamics is an important consequence of this research. Considering Reichenbach's theorem, we can deduce that either the scoreline causes information dynamics, or the information dynamics causes the scoreline, or there is a common cause for the two measures of performance. The first two cases are unlikely, and thus we conjecture that the information dynamics and scoreline are proxy to an underlying cause. Further, our results support the hypothesis of intrinsic motivation in psychology and reinforcement learning [9], whereby it is shown that an embodied agent that is both intrinsically and extrinsically motivated

is more adept at problem solving. In the case of team dynamics, information dynamics is an intrinsic reward and scoring goals is an extrinsic reward. This relates to the work of Zahedi et al. [72], who used a linear combination of predictive information to speed up the learning process of an embodied agent.


6.3 State Space Coherence Diagrams

Figure 6 shows tactical information dynamics, that is, state-space coherence diagrams for all tactical roles, while Figures 7 and 8 show partitioned information dynamics: state space coherence diagrams for specific tactical roles, color-coded with the scorelines.

Both state space coherence diagrams in Figure 6 clearly show separation among three tactical roles: defenders, midfielders, and attackers. Each tactical role is clustered well in each of the contexts. Defenders (shown in red) tend to have low relative rigidity and low relative responsiveness.



(a) $\delta A(\mathcal{G}, \mathcal{C})$ versus $\delta T(\mathcal{G}, \mathcal{C})$



(b) $\delta A(\mathcal{G}, \mathcal{H})$ versus $\delta T(\mathcal{G}, \mathcal{H})$

Figure 6. Tactical information dynamics: state space coherence diagrams of relative responsiveness $\delta T(\mathcal{G}, \mathcal{Y})$ as a function of relative rigidity $\delta A(\mathcal{G}, \mathcal{Y})$, with two opponents of team Gliders (\mathcal{G}): team Cyrus (\mathcal{C}), and team HELIOS (\mathcal{H}). In the diagrams, red points are used for relative responsiveness versus relative rigidity with respect to Gliders' defenders, with green points for Gliders' midfielders and blue for attackers.

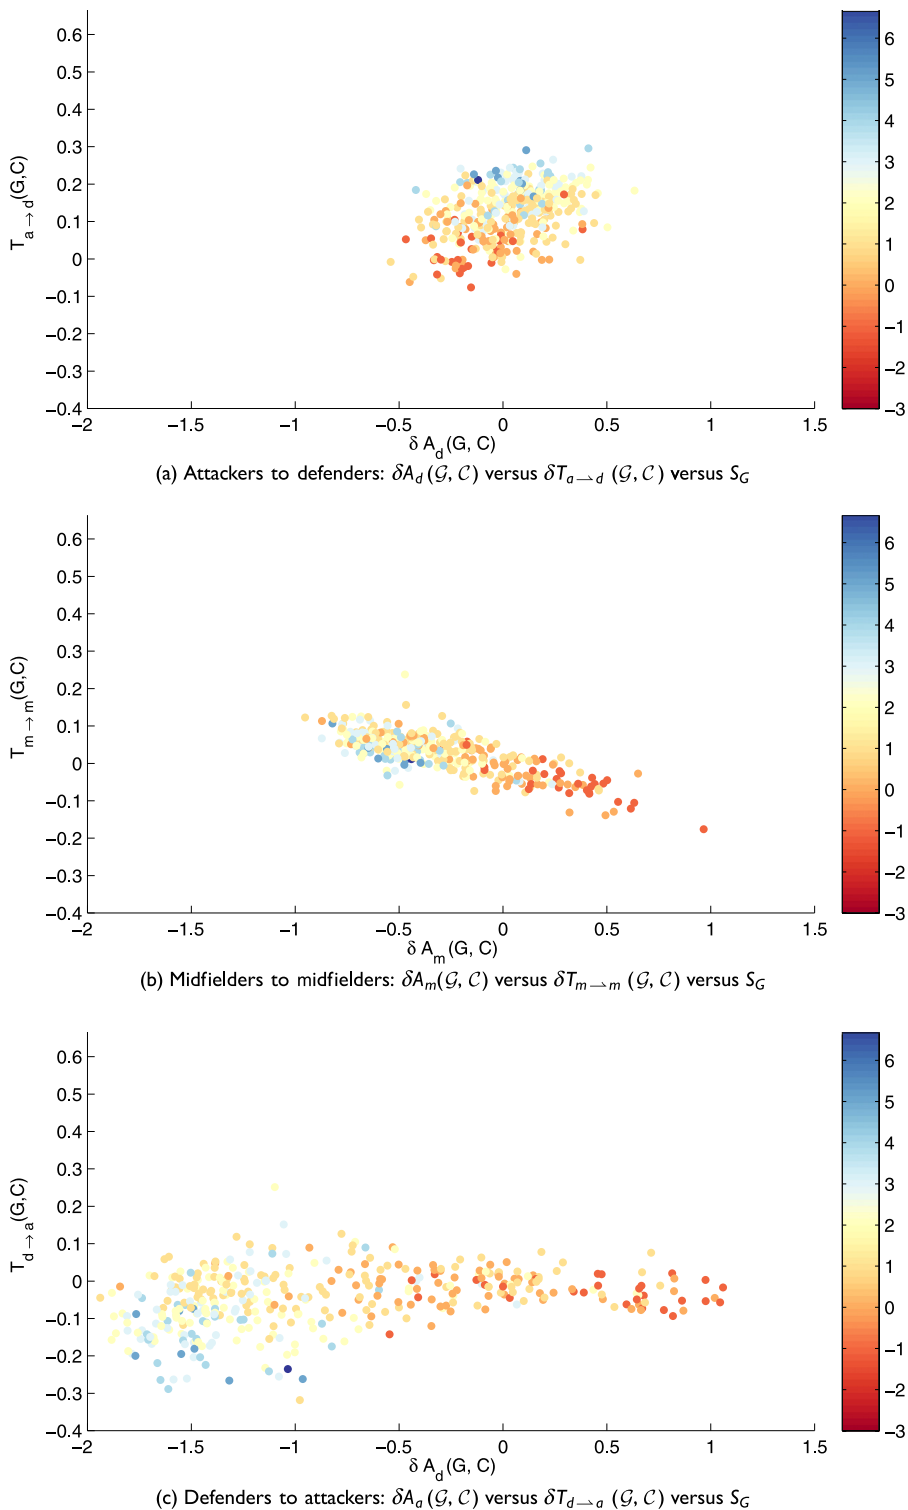


Figure 7. Partitioned information dynamics: state-space coherence diagrams for specific tactical roles, with colormapping of the correlation with scorelines. Relative responsiveness $\delta T(\mathcal{G}, \mathcal{C})$ is a function of relative rigidity $\delta A(\mathcal{G}, \mathcal{C})$ for Gliders (\mathcal{G}) versus Cyrus (\mathcal{C}).

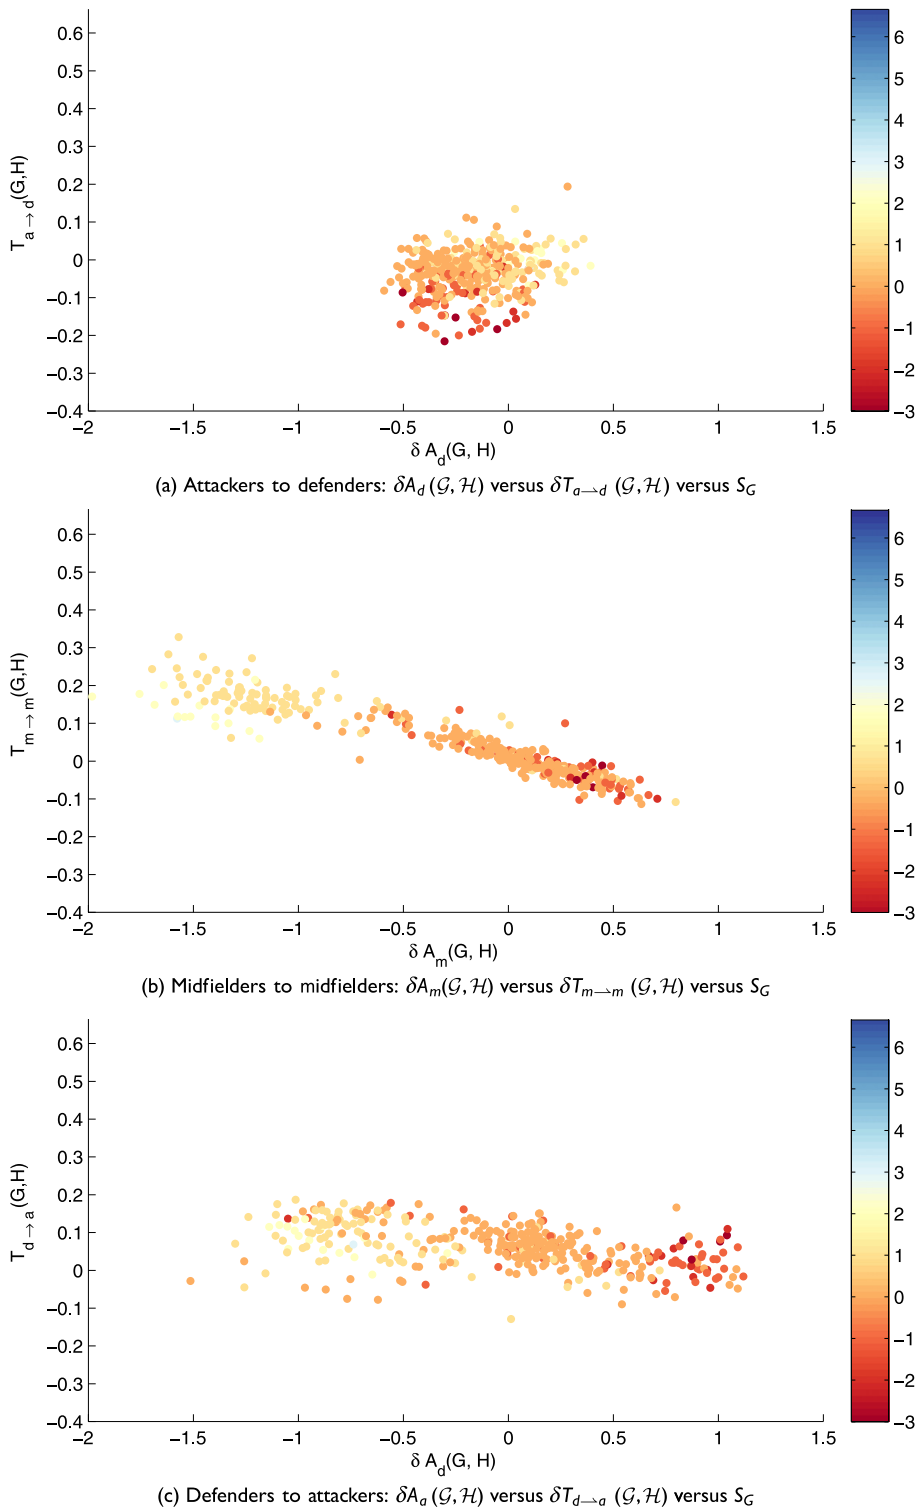


Figure 8. Partitioned information dynamics: state space coherence diagrams for specific tactical roles, with colormapping of the correlation with scorelines. Relative responsiveness $\delta T(\mathcal{G}, \mathcal{H})$ is a function of relative rigidity $\delta A(\mathcal{G}, \mathcal{H})$ for Gliders (\mathcal{G}) versus HELIOS (\mathcal{H}).

That is, defenders of the competing teams in each contest (Gliders versus Cyrus and Gliders versus HELIOS) do not differ much in their rigidity and responsiveness, except that Gliders defenders are more responsive than Cyrus defenders. Midfielders (shown in green) consistently occupy a well-defined narrow region, showing that an increase in relative rigidity is correlated with a decrease in relative responsiveness, in both contests. Gliders midfielders appear to be slightly more responsive and less rigid than HELIOS midfielders. Finally, attackers (shown in blue) are clustered differently in two contests. In games between Gliders and Cyrus, low relative rigidity is correlated with a wider spread of relative responsiveness, which tends to be negative. In other words, when Gliders attackers are less rigid than Cyrus attackers, they are also less responsive: This is indicative of their more explorative behavior around and within their opponent's penalty area. This feature is not observed in the diagram for Gliders versus HELIOS; moreover, there is a correlation between relative rigidity and responsiveness similar to the one in the midfielders' cluster. This reinforces an earlier observation that in the games between Gliders and HELIOS, the attackers often play in the midfield.

Importantly, these state space coherence diagrams allow us to examine average role-based multi-agent dynamics across games, by clustering dynamic processes in an abstract state space and identifying salient features of competing tactical formations.

Now we turn our attention to information dynamics partitioned for each tactical role and their correlation with the scorelines. The partitioned diagrams in Figures 7 and 8 reveal how the differences in rigidity and responsiveness are consistently related to the performance, across both contests. For example, there is a clear correlation between better performance and higher responsiveness and higher rigidity of defenders, as shown in Figures 7a and 8a. As mentioned earlier, a positive contribution of the higher rigidity is not counterintuitive, as it results from synchronous, and hence more predictable on average, movement of each defender following the line defense tactic, enabling efficient offside traps for the opposition. On the other hand, for the midfielders, there is a clear correlation between better scorelines and lower rigidity as well as higher responsiveness, as shown in Figures 7b and 8b. That is, when Gliders midfielders are less rigid or more responsive than the opponent's midfielders, the Gliders team tends to win. Finally, it is evident that when Gliders attackers are less rigid and less responsive than Cyrus attackers (Figure 7c), the team benefits, while in the games versus HELIOS the correlation with performance is mostly observed for lower rigidity (Figure 8c). The difference between two contests is again due to the fact that Gliders attackers are typically restrained to playing in midfield in the games versus HELIOS.

These partitioned diagrams provide another useful tool in clustering the multi-agent dynamics and classifying the games in terms of tactical behavior.

7 Conclusion

In this article we have addressed two problems: (i) identifying interaction networks that link together autonomous agents, using only the observational data without reconstructing the agents' control logic and internal behavior; and (ii) classifying coherent dynamic situations within the multi-agent games, in the context of distributed communications. The methodology is not aimed at explicit interactions within a team, but rather at implicit interactions, across teams, that may be spatially long-range. The approach to constructing interaction networks used a novel application of information dynamics analyzing pairwise interactions and role-based tactics, exemplified by RoboCup 2D Simulation League games.

The interaction networks were demonstrated with two network subtypes: information-sink and information-source diagrams. In an information-sink diagram every node (every player) has an incoming edge, while in an information-source diagram every node has an outgoing edge. These diagrams represent simplifications to full effective network diagrams, and while they do not reveal the full interaction structure, they are significantly more efficient to compute, and highlight the strongest of the interactions. Information-sink and -source diagrams were computed for two experimental setups that matched the RoboCup-2014 vice-champion team Gliders [53] against two top-five

teams, Cyrus [24] and HELIOS [3], and showed, for the first time, a number of asymmetries in the tactical schemes used by the teams. These quantified asymmetries were used in allocating suitable defensive resources by team Gliders, resulting in a statistically significant performance gain.

The follow-up analysis involved computation of information transfer and storage used to quantify (relative) responsiveness and rigidity, respectively. These notions can be applied to individual agents, tactical roles of agents, and the team overall. Both measures, relative responsiveness and rigidity, were correlated with the game results, pointing out important couplings in particularly intense interactions across teams, and highlighting the tactical roles and field areas where the game outcomes were mostly decided.

We then examined average role-based multi-agent dynamics across games via novel state space coherence diagrams, which clustered the dynamic processes in an abstract state space. In our examples, the state space plots identified several salient features of competing tactical formations, providing a crucial step in classifying the games in terms of tactical behavior. In general, these diagrams are useful when there is a need to cluster dynamic, rather than static, processes.

The information dynamics tools introduced in this article are applicable in several artificial life and biological scenarios, where an accurate estimation of the information-processing channels can reveal a computational structure underlying the emergence of collective behaviors.

Several related simulation-based studies are worth pointing out that used information dynamics to identify leadership within a swarm, for example, leaderships in pairs of zebrafish [7], and covert leadership in a swarm of robots distinguished by transfer entropy [63]. While a leader is defined as a swarm member that acts upon specific information in addition to what is provided by local interactions [12, 62], a covert leader is treated no differently than others in the swarm, so that leaders and followers interact identically [55]. By contrasting transfer entropies across individuals, the study [63] was able to distinguish the covert leaders from the followers by characterizing the covert leaders with a lesser amount of transfer entropy than the followers. Furthermore, perhaps counterintuitively, the leaders do not share more information with the swarm than the followers do. In the context of this article, the followers may be seen as larger information sinks than the leaders, highlighting another potential use of the information-sink diagrams.

Similar information dynamics measures have also been very recently used to measure pairwise correlations in a biological swarm of soldier crabs [65], revealing that in smaller swarms the crabs tend to make decisions based on their own past behavior, whereas in larger swarms they make decisions based on the behavior of their neighbors rather than their own.

One possible direction of future research is to investigate how each tactical role could correspond to a different relation between rigidity and responsiveness, and relate these to components of the information-theoretic measure of autonomy [5]. Ultimately, the analysis can be extended to include comprehensive tactical planning and decision making.

We believe that the methodology proposed in this article would be useful not only in multi-agent team sports, but also in various analyses of distributed dynamics, such as decentralized coordination [21, 71], swarm engineering [41, 68], and evolutionary robotics [19, 46, 47, 60].

Acknowledgments

This analysis was carried out using team Gliders-2014 [2, 48, 50, 53]—a vice-champion team of the RoboCup Soccer Simulation League 2D. Further, the team Gliders-2016 [49] became the world champions of RoboCup-2016.

References

1. Abreu, P. H., Moura, J., Silva, D. C., Reis, L. P., & Garganta, J. (2012). Performance analysis in soccer: A cartesian coordinates based approach using RoboCup data. *Soft Computing*, 16, 47–61.
2. Akiyama, H. (2010). Agent2d base code. <http://www.rctools.sourceforge.jp>.
3. Akiyama, H., Nakashima, T., Yamashita, K., & Mifune, S. (2014). *HELIOS2014 team description paper* (Technical report). RoboCup 2014 Symposium and Competitions: Team Description Papers.

4. Asada, M., Kitano, H., Noda, I., & Veloso, M. (1999). RoboCup: Today and tomorrow—what we have learned. *Artificial Intelligence*, *110*, 193–214.
5. Bertschinger, N., Olbrich, E., Ay, N., & Jost, J. (2008). Autonomy: An information-theoretic perspective. *Biosystems*, *91*, 331–345.
6. Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., & Walczak, A. M. (2012). Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences of the U.S.A.*, *109*, 4786–4791.
7. Butail, S., Mwaffo, V., & Porfiri, M. (2016). Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Physical Review E*, *93*, 042411.
8. Butler, M., Prokopenko, M., & Howard, T. (2001). Flexible synchronisation within RoboCup environment: A comparative analysis. In P. Stone, T. Balch, & G. Kraetzschmar (Eds.), *RoboCup 2000: Robot Soccer World Cup IV* (pp. 119–128). Berlin, Heidelberg: Springer.
9. Chentanez, N., Barto, A. G., & Singh, S. P. (2004). Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems* (pp. 1281–1288). Cambridge, MA: MIT Press.
10. Cliff, O. M., Lizier, J. T., Wang, X. R., Wang, P., Obst, O., & Prokopenko, M. (2014). Towards quantifying interaction networks in a football match. In S. Behnke, M. Veloso, A. Visser, & R. Xiong (Eds.), *RoboCup 2013: Robot World Cup XVII* (pp. 1–12). Berlin, Heidelberg: Springer.
11. Cliff, O. M., Prokopenko, M., & Fitch, R. (2016). An information criterion for inferring coupling in distributed dynamical systems. *Frontiers in Robotics & AI*, *3*.
12. Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, *433*, 513–516.
13. Crutchfield, J. P., & Feldman, D. P. (2003). Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos*, *13*, 25–54.
14. Duch, J., Waitzman, J. S., & Amaral, L. A. N. (2010). Quantifying the performance of individual players in a team activity. *PLOS ONE*, *5*, e10937.
15. Faes, L., Nollo, G., & Porta, A. (2011). Information-based detection of nonlinear Granger causality in multivariate processes via a nonuniform embedding technique. *Physical Review E*, *83*, 051112.
16. Faes, L., Nollo, G., & Porta, A. (2012). Non-uniform multivariate embedding to assess the information transfer in cardiovascular and cardiorespiratory variability series. *Computers in Biology and Medicine*, *42*, 290–297.
17. Feldman, D. P., McTague, C. S., & Crutchfield, J. P. (2008). The organization of intrinsic computation: Complexity-entropy diagrams and the diversity of natural information processing. *Chaos*, *18*, 043106.
18. Fewell, J., Armbruster, D., Ingraham, J., Petersen, A., & Waters, J. (2012). Basketball teams as strategic networks. *PLOS ONE*, *7*, e47445.
19. Fitch, R., & Butler, Z. J. (2008). Million module march: Scalable locomotion for large self-reconfiguring robots. *The International Journal of Robotics Research*, *27*, 331–343.
20. Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, *2*, 56–78.
21. Gan, S., Fitch, R., & Sukkarieh, S. (2014). Online decentralized information gathering with spatial–temporal constraints. *Autonomous Robots*, *37*, 1–25.
22. Gómez-Herrero, G., Wu, W., Rutanen, K., Soriano, M., Pipa, G., & Vicente, R. (2015). Assessing coupling dynamics from an ensemble of time series. *Entropy*, *17*, 1958–1970.
23. Kassir, A., Fitch, R., & Sukkarieh, S. (2015). Communication-aware information gathering with dynamic information flow. *The International Journal of Robotics Research*, *34*, 173–200.
24. Khayami, R., Zare, N., Karimi, M., Mahor, P., Afshar, A., Najafi, M. S., Asadi, M., Tekrar, F., Asali, E., & Keshavarzi, A. (2014). *CYRUS 2D simulation team description paper 2014* (Technical report). RoboCup 2014 Symposium and Competitions: Team Description Papers.
25. Kijima, A., Yokoyama, K., Shima, H., & Yamamoto, Y. (2014). Emergence of self-similarity in football dynamics. *The European Physical Journal B*, *87*, 1–6.

26. Kitano, H., & Asada, M. (1998). The RoboCup humanoid challenge as the millennium challenge for advanced robotics. *Advanced Robotics*, *13*, 723–736.
27. Kitano, H., Tambe, M., Stone, P., Veloso, M., Coradeschi, S., Osawa, E., Matsubara, H., Noda, I., & Asada, M. (1998). The RoboCup synthetic agent challenge 97. In H. Kitano (Ed.), *RoboCup 1997: Robot Soccer World Cup I* (pp. 62–73). Berlin, Heidelberg: Springer.
28. Kok, J. R., Spaan, M. T. J., & Vlassis, N. A. (2005). Non-communicative multi-robot coordination in dynamic environments. *Robotics and Autonomous Systems*, *50*, 99–114.
29. Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. *Physical Review E*, *69*, 066138.
30. Liu, W., & Winfield, A. (2010). Modelling and optimisation of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, *29*, 1741–1760.
31. Lizier, J. T. (2014). JIDT: An information-theoretic toolkit for studying the dynamics of complex systems. *Frontiers in Robotics and AI*, *1*, 11.
32. Lizier, J. T. (2014). Measuring the dynamics of information processing on a local scale in time and space. In M. Wibral, R. Vicente, & J. T. Lizier (Eds.), *Directed information measures in neuroscience* (pp. 161–193). Berlin, Heidelberg: Springer.
33. Lizier, J. T., Heinzele, J., Horstmann, A., Haynes, J.-D., & Prokopenko, M. (2011). Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *Journal of Computational Neuroscience*, *30*, 85–107.
34. Lizier, J. T., Prokopenko, M., & Zomaya, A. Y. (2008). Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, *77*, 026110.
35. Lizier, J. T., Prokopenko, M., & Zomaya, A. Y. (2010). Information modification and particle collisions in distributed computation. *Chaos*, *20*, 037109.
36. Lizier, J. T., Prokopenko, M., & Zomaya, A. Y. (2012). Coherent information structure in complex computation. *Theory in Biosciences*, *131*, 193–203.
37. Lizier, J. T., Prokopenko, M., & Zomaya, A. Y. (2012). Local measures of information storage in complex distributed computation. *Information Sciences*, *208*, 39–54.
38. Lizier, J. T., & Rubinov, M. (2012). *Multivariate construction of effective computational networks from observational data* (Technical report). Max Planck Institute for Mathematics in the Sciences.
39. Marinazzo, D., Pellicoro, M., & Stramaglia, S. (2012). Causal information approach to partial conditioning in multivariate data sets. *Computational and Mathematical Methods in Medicine*, *2012*, 303601+.
40. Mathews, G., Durrant-Whyte, H. F., & Prokopenko, M. (2009). Decentralised decision making in heterogeneous teams using anonymous optimisation. *Robotics and Autonomous Systems*, *57*, 310–320.
41. Miller, J. M., Wang, X. R., Lizier, J. T., Prokopenko, M., & Rossi, L. F. (2014). Measuring information dynamics in swarms. In M. Prokopenko (Ed.), *Guided self-organization: Inception* (pp. 343–364). Berlin, Heidelberg: Springer.
42. Mota, L., Reis, L. P., & Lau, N. (2011). Multi-robot coordination using setplays in the middle-size and simulation leagues. *Mechatronics*, *21*, 434–444.
43. Noda, I., & Stone, P. (2003). The RoboCup soccer server and CMUnited clients: Implemented infrastructure for MAS research. *Autonomous Agents and Multi-Agent Systems*, *7*, 101–120.
44. Peña, J. L., & Touchette, H. (2013). A network theory analysis of football strategies. In C. Clanet (Ed.), *Proceedings of Euromech Physics of Sports Conference* (pp. 517–528). Éditions de l'École Polytechnique.
45. Pitonakova, L., Crowder, R., & Bullock, S. (2016). Information flow principles for plasticity in foraging robot swarms. *Swarm Intelligence*, *10*, 33–63.
46. Prokopenko, M., Gerasimov, V., & Tanev, I. (2006). Evolving spatiotemporal coordination in a modular robotic system. In S. Nolfi, G. Baldassarre, R. Calabretta, J. C. T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, & D. Parisi (Eds.), *From animals to animats* (pp. 558–569). Berlin, Heidelberg: Springer.
47. Prokopenko, M., Gerasimov, V., & Tanev, I. (2006). Measuring spatiotemporal coordination in a modular robotic system. In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone, & A. Vespignani (Eds.), *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems* (pp. 185–191). Cambridge, MA: MIT Press.

48. Prokopenko, M., Obst, O., Wang, P., Budden, D., & Cliff, O. (2013). *Gliders2013: Tactical analysis with information dynamics* (Technical report). RoboCup 2013 Symposium and Competitions: Team Description Papers.
49. Prokopenko, M., Wang, P., Obst, O., & Jauregui, V. (2016). *Gliders2016: Integrating multi-agent approaches to tactical diversity* (Technical report). RoboCup 2016 Symposium and Competitions: Team Description Papers.
50. Prokopenko, M., Obst, O., Wang, P., & Held, J. (2012). *Gliders2012: Tactics with action-dependent evaluation functions* (Technical report). RoboCup 2012 Symposium and Competitions: Team Description Papers.
51. Prokopenko, M., & Wang, P. (2003). Evaluating team performance at the edge of chaos. In D. Polani, B. Browning, A. Bonarini, & K. Yoshida (Eds.), *RoboCup 2003: Robot Soccer World Cup VII* (pp. 89–101). Berlin, Heidelberg: Springer.
52. Prokopenko, M., & Wang, P. (2003). Relating the entropy of joint beliefs to multi-agent coordination. In G. A. Kaminka, P. U. Lima, & R. Rojas (Eds.), *RoboCup 2002: Robot Soccer World Cup VI* (pp. 367–374). Berlin, Heidelberg: Springer.
53. Prokopenko, M., Wang, P., & Obst, O. (2014). *Gliders2014: Dynamic tactics with Voronoi diagrams* (Technical report). RoboCup 2014 Symposium and Competitions: Team Description Papers.
54. Riedmiller, M., Gabel, T., Hafner, R., & Lange, S. (2009). Reinforcement learning for robot soccer. *Autonomous Robots*, 27, 55–73.
55. Rossi, L., Han, X., & Shen, C.-C. (2007). Autonomous navigation of wireless robot swarms with covert leaders. In *Proceedings of the International Conference on Robot Communication and Coordination*. ICST.
56. Schreiber, T. (2000). Measuring information transfer. *Physical Review Letters*, 85, 461–464.
57. Sporns, O. (2011). *Networks of the brain*. Cambridge, MA: MIT Press.
58. Stetter, O., Battaglia, D., Soriano, J., & Geisel, T. (2012). Model-free reconstruction of excitatory neuronal connectivity from calcium imaging signals. *PLOS Computational Biology*, 8, e1002653.
59. Stone, P., & Veloso, M. (1999). Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110, 241–273.
60. Stoy, R. F., Kernbach, S., Nagpal, R., & Shen, W. (2014). Special issue: Reconfigurable modular robotics. *Robotics and Autonomous Systems*, 62, 943–1084.
61. Stramaglia, S., Wu, G.-R., Pellicoro, M., & Marinazzo, D. (2012). Expanding the transfer entropy to identify information circuits in complex systems. *Physical Review E*, 86, e066211.
62. Sun, Y., Rossi, L. F., Luan, H., & Shen, C.-C. (2013). Modeling and analyzing large swarms with covert leaders. In M. Parashar & F. Zambonelli (Eds.), *Proceedings of the International Conference on Self-Adaptive and Self-Organizing Systems* (pp. 169–178). New York: IEEE.
63. Sun, Y., Rossi, L. F., Shen, C.-C., Miller, J., Wang, X. R., Lizier, J. T., Prokopenko, M., & Senanayake, U. (2014). Information transfer in swarms with leaders. In J. Nickerson & T. Malone (Eds.), *Proceedings of the Collective Intelligence Conference*, arXiv.org.
64. Taylor, M. E., & Stone, P. (2009). Transfer learning for reinforcement learning domains: A survey. *The Journal of Machine Learning Research*, 10, 1633–1685.
65. Tomaru, T., Murakami, H., Niizato, T., Nishiyama, Y., Sonoda, K., Moriyama, T., & Gunji, Y.-P. (2016). Information transfer in a swarm of soldier crabs. *Artificial Life and Robotics*, 21, 177–180.
66. Vicente, R., Wibral, M., Lindner, M., & Pipa, G. (2011). Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30, 45–67.
67. Vilar, L., Araujo, D., Davids, K., & Bar-Yam, Y. (2013). Science of winning soccer: Emergent pattern-forming dynamics in association football. *Journal of Systems Science and Complexity*, 26, 73–84.
68. Wang, X. R., Miller, J. M., Lizier, J. T., Prokopenko, M., & Rossi, L. F. (2012). Quantifying and tracing information cascades in swarms. *PLOS ONE*, 7, e40084.
69. Whiteson, S., Kohl, N., Miikkulainen, R., & Stone, P. (2005). Evolving keepaway soccer players through task decomposition. *Machine Learning*, 59, 5–30.

70. Wibral, M., Rahm, B., Rieder, M., Lindner, M., Vicente, R., & Kaiser, J. (2011). Transfer entropy in magnetoencephalographic data: Quantifying information flow in cortical and cerebellar networks. *Progress in Biophysics and Molecular Biology*, *105*, 80–97.
71. Xu, Z., Fitch, R., Underwood, J. P., & Sukkarieh, S. (2013). Decentralized coordinated tracking with mixed discrete-continuous decisions. *Journal of Field Robotics*, *30*, 717–740.
72. Zahedi, K., Martius, G., & Ay, N. (2013). Linear combination of one-step predictive information with an external reward in an episodic policy gradient setting: A critical analysis. *Frontiers in Psychology*, *131*.