Gliders2013: Tactical Analysis with Information Dynamics

Mikhail Prokopenko, Oliver Obst, Peter Wang, David Budden, and Oliver Cliff

CSIRO Information and Communication Technologies Centre, Adaptive Systems,
PO Box 76, Epping, NSW 1710, Australia

Abstract. The RoboCup 2D Simulation League incorporates several challenging features, setting a benchmark for Artificial Intelligence. In this paper we describe some of the ideas and tools used in development of our team, Gliders2013. In doing so, we focus on information dynamics as one of the central mechanisms for tactical analysis. This analysis involves computation of information transfer and storage, relating the information transfer to responsiveness of the players, and the information storage within the team to the team's rigidity and lack of tactical richness. The proposed approach has been successfully applied to tactical opponent modelling.

1 Introduction

Autonomous decision-making, under constraints enforced by tactical plans and teamwork (collaboration) as well as opponent (competition) is one of the most prominent characteristics of the RoboCup 2D Simulation League [1,2,3,4,5,6,7,8].

Information dynamics is a recent methodology for analysis of complex systems in general and swarm behavior in particular. In this short paper we describe a novel application of information dynamics to the tactical analysis of RoboCup 2D Simulation League games. This analysis was carried out using Gliders2013, a follow up on Gliders2012 [8] — a simulated soccer team for the RoboCup soccer 2D simulator [9]. Gliders2012 participated in RoboCup 2012 with convincing results. The tournament included 19 teams and was played over several rounds after which Gliders2012 proceeded to the semi-finals, resulting in a 4th place. The team code is written in C++ using agent2d: the well-known base code developed by Akiyama et al. [10]. Other software packages are used as well:

- librcsc: a base library for RCSS with various utilities describing relevant geometrical constructs, world model, etc.;
- soccerwindow2: a viewer program for RCSS, working as a monitor client, a log player and a visual debugger;
- fedit2: a formation editor for agent2d, allowing to design a team formation;
- Gliders' in-browser basic soccermonitor (GIBBS): a log-player for viewing 2D Simulation League logs over web browser [11];
- we are exploring options to use released source code of some RoboCup 2012 teams, e.g., Marlik [12], robOTTO [13], etc.

GIBBS is used to visualize matches in a browser window [11]. At this time, the main implementation uses HTML5 / javascript, with a pre-processing step to create simple-to-parse log files that can be conveniently transferred over the Internet. The current state of the visualization is basic but functional, and it is released by us to the RoboCup community in the hope that it will be improved over time, re-released, and be used for showing games to general public in future competitions. The GIBBS log player is available (along with the necessary converter) at the Gliders2012 web page www.oliverobst.eu/research/robotics-gliders2012-simulation-league-robocup-team.

2 Motivation and approach

To establish a recent baseline, we carried out multiple iterative experiments matching Gliders2012 up against teams not based on agent2d. Among the latter class, we chose Oxsy [14], achieving $\approx +0.05$ goal difference, typically averaged over 500 games. This baseline was improved to $\approx +0.50$ goal difference in the team Gliders2013 that had its tactics adjusted as a result of tactical analysis described in the following section.

2.1 Information dynamics

A recently developed framework of *information dynamics* studies the phenomenon of computation in a systematic way: it uncovers and analyses information-theoretic roots of the most basic computational primitives: *storage*, *transmission*, and *modification of information* [15,16,17,18,19].

The *active information storage* quantifies the information storage component that is directly in use in the computation of the next state of a process [15,19]. More precisely, it is the average mutual information between the semi-infinite past of the process $x_n^{(k)} = \{x_{n-k+1}, \ldots, x_{n-1}, x_n\}$ (as $k \to \infty$) and its next state. The *local information storage* (or pointwise mutual information) is then a measure of the amount of information storage in use by the process at a particular time-step $n+1$:

$$a_X(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_n^{(k)}, x_{n+1})}{p(x_n^{(k)})p(x_{n+1})}. \tag{1}$$

For example, the local active information storage in Cellular Automata (CA) provides evidence that periodic nonmoving structures (e.g., blinkers in Conway's Game of Life) and background domains are dominant information storage processes in these systems. In practice, one deals with finite-k estimates $a_X(n+1,k)$, as well as the finite-k estimates $A_X(k)$ of the average active information storage $A_X = \langle a_X(n+1)\rangle_n$.

Transfer entropy [20] is designed to detect asymmetry in the interaction of subsystems by distinguishing between "driving" and "responding" elements. The *local information transfer*, based on *transfer entropy*, captures information transmission [16] from source Y to destination X, at a particular time-step $n+1$. Specifically, the local information transfer between a source and a destination agent is defined as the information provided by the source y_n about the destination's next state x_{n+1} that was not contained in the past of the destination $x_n^{(k)}$:

$$t_{Y\to X}(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n)}{p(x_{n+1} \mid x_n^{(k)})}. \tag{2}$$

When applied to CA for filtering coherent structures it provides the quantitative evidence for the long-held conjecture that domain walls and traveling coherent structures (e.g., gliders in Conway's Game of Life) are the dominant information transfer agents in CA. This method is distinguished in using asymmetric, multivariate, information-theoretical analysis, which captures not only directional and non-linear relationships, but also collective interactions. Interestingly, the local information transfer can be negative, indicating that the source is misinformative about the next state of the destination.

It is important to realise that information transfer between two variables does not require an eplicit communication channel, it rather indicates a high degree of directional synchrony or nonlinear correlation between the source and the destination. It characterises a degree of *predictive* information transfer, i.e., "if the state of the source is known, how much does that help to predict the state of the destination?" [16].

Sometimes it is useful to condition the local information transfer on another contributing process W, considering the *local conditional transfer entropy* [17]:

$$t_{Y\to X|W}(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n, w_n)}{p(x_{n+1} \mid x_n^{(k)}, w_n)}. \tag{3}$$

2.2 Tactical analysis

Building up on the information dynamics measures, it is possible to investigate group behavior in complex systems, such as swarms. For instance, recent studies by Wang et al. [21] 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. The study of Wang et al. [21] proposed and verified another hypothesis according to which information cascades that correspond to long range communications are captured by *conditional transfer entropy* [16,17]. In other words, information transfer was shown to characterise the communication aspect of collective computation distributed within the swarm.

In applying information dynamics to the RoboCup 2D Simulation League we make the following conjecture:

> *a higher information transfer $t_{Y\to X|W}$ 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 W (e.g., the movement of the ball W), is indicative of a higher responsiveness of the process/player X to the process/player Y.*

That is, the "destination" player Y responds, for example, by repositioning, to the 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 opponent's defenders we consider the information transfer $t_{Y_{def}\to X_{att}}$ from defenders $y_i \in Y_{def}$ to forwards $x_j \in X_{att}$, where the involved probability distributions are obtained for different relative positions on the soccer field. Vice versa, dynamics of the opponent's defenders trying to better mark our team's forwards are represented in the information transfer $t_{X_{att}\to Y_{def}}$ from forwards $x_j \in X_{att}$ to defenders $y_i \in Y_{def}$. These two examples specifically consider a coupling between the attack line X_{att} of our team and the defense line Y_{def} of opponent's team (henceforth we keep denoting opponent's lines by Y and our team's lines by X).

One may further contrast these two transfers in the coupled lines:

$$\Delta(X_{att}, Y_{def}) = t_{Y_{def} \to X_{att}} - t_{X_{att} \to Y_{def}} \tag{4}$$

When our forwards are more responsive on average to the opponent's defenders than the opponents defenders are to our forwards, $t_{Y_{def} \to X_{att}} > t_{X_{att} \to Y_{def}}$, and the relative responsiveness $\Delta(X_{att}, Y_{def}) > 0$.

It is clear that one may consider various coupled lines (our attack vs opponent's defense; our defense vs opponent's attack; midfield vs midfield, etc.), dependent on the deployed formations. In is also possible to combine relative responsiveness scores for each of the coupled lines in an overall tactical relative responsiveness (including, for example, relative scores for midfielders X_{mid} and Y_{mid}):

$$\begin{aligned} \Delta(X_{team}, Y_{team}) &= t_{Y_{def} \to X_{att}} - t_{X_{att} \to Y_{def}} \quad (5)\\ &+ t_{Y_{att} \to X_{def}} - t_{X_{def} \to Y_{att}} \quad (6)\\ &+ t_{Y_{mid} \to X_{mid}} - t_{X_{mid} \to Y_{mid}} \quad (7) \end{aligned}$$

We would like to point out that all the transfers to team X are added up, and the transfers from team X are substracted away. When each of the transfers is conditioned on some othwer contributor W (for example, all the dynamics are computed in the context of the ball movement), the overall tactical relative responsiveness $\Delta(X_{team}, Y_{team}|W)$ is also placed in this specific context of the game, W.

In principle, competitive situations result in quite vigorous dynamics, repositioning and responses within the involved lines and overall formations, and the team that manages to achieve a higher degree of tactical relative responsiveness does in general perform better. Ultimately, one may correlate the scores of relative responsiveness (either line-by-line or overall) with the game scores, and identify the lines which impacted on the games more. For example, after analysing the games between Gliders and Oxsy, we have observed that high correlation exists between the game score and only one relative responsiveness $\Delta(X_{att}, Y_{def})$. That is, the games between these two teams are decided mostly in the opposition between Gliders' forwards and Oxsy's defenders: whenever the forwards are more responsive in evading the marks, Gliders tend to win, and whenever Oxsy's defenders are more agile to closing on to the forwards, Gliders tend to lose. On the contrary, the contests between Gliders' defenders and Oxsy's forwards, or the dynamic encounters between the midfield players, do not seem to affect the scoreline on average.

Our tactical analysis also involves computation of the active information storage within the teams hypothesised to characterise team's rigidity A_X, followed by determining the relative rigidity $A(X, Y) = A_X - A_Y$ for the teams (or their coupled lines). The intuition here is that if the rigidity A_X is high then this is counter-productive, since the players are not as independent of each other as a truly complex or swarm behavior would warrant. Consequently, the relative rigidity $A(X, Y)$ should be anti-correlated with the team X performance against team Y.

A more detailed analysis is being finalised.

3 Conclusion

We described a novel application of information dynamics to the tactical analysis of RoboCup 2D Simulation League games. This analysis involved computation of information transfer and storage, and two hypotheses. The first one related positive information transfer from players Y to players X as an indication of responsiveness of the latter, suggesting to compute relative responsiveness between the opposing lines of two teams. The second hypothesis connected the information storage within the team with the team's rigidity, harming the fluidity and tactical richness of the team. This relation yielded the score for relative rigidity between the opposing teams. Both measures, relative responsiveness and rigidity, were correlated with the game results, and the preliminary observations support the hypotheses.

The proposed approach has been successfully applied to opponent modelling and selecting the best available tactics in an opponent-specific way.

Acknowledgments The Authors are thankful to Ivan Duong, Edward Moore and Jason Held for their contribution to Gliders2012 and GIBBS, as well as Joseph T. Lizier and Rosalind X. Wang for their help with developing and applying information dynamcis measures. Team logo was created by Matthew Chadwick. Some of the Authors have been involved with RoboCup Simulation 2D in the past, however the code of their previous teams (Cyberoos and RoboLog, see, e.g., [7,22]) is not used in Gliders.

References

1. Noda, I., Stone, P.: The RoboCup Soccer Server and CMUnited Clients: Implemented Infrastructure for MAS Research. Autonomous Agents and Multi-Agent Systems **7**(1–2) (July–September 2003) 101–120
2. Riley, P., Stone, P., Veloso, M.: 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, Berlin (2001)
3. Stone, P., Riley, P., Veloso, M.: Defining and using ideal teammate and opponent models. In: Proceedings of the Twelfth Annual Conference on Innovative Applications of Artificial Intelligence. (2000)
4. Butler, M., Prokopenko, M., Howard, T.: Flexible synchronisation within RoboCup environment: A comparative analysis. In: RoboCup 2000: Robot Soccer World Cup IV, London, UK, Springer (2001) 119–128
5. Reis, L.P., Lau, N., Oliveira, E.: Situation based strategic positioning for coordinating a team of homogeneous agents. In: Balancing Reactivity and Social Deliberation in Multi-Agent Systems, From RoboCup to Real-World Applications (selected papers from the ECAI 2000 Workshop and additional contributions), London, UK, Springer (2001) 175–197
6. Prokopenko, M., Wang, P.: Relating the entropy of joint beliefs to multi-agent coordination. In Kaminka, G.A., Lima, P.U., Rojas, R., eds.: RoboCup 2002: Robot Soccer World Cup VI. Volume 2752 of Lecture Notes in Computer Science., Springer (2003) 367–374
7. Prokopenko, M., Wang, P.: 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. Volume 3020 of Lecture Notes in Computer Science., Springer (2003) 89–101

8. Prokopenko, M., Obst, O., Wang, P., Held, J.: Gliders2012: Tactics with action-dependent evaluation functions. In: RoboCup 2012 Symposium and Competitions: Team Description Papers, Mexico City, Mexico, June 2012. (2012)
9. Chen, M., Dorer, K., Foroughi, E., Heintz, F., Huang, Z., Kapetanakis, S., Kostiadis, K., Kummeneje, J., Murray, J., Noda, I., Obst, O., Riley, P., Steffens, T., Wang, Y., Yin, X.: Users Manual: RoboCup Soccer Server — for Soccer Server Version 7.07 and Later. The RoboCup Federation. (February 2003)
10. Akiyama, H.: Agent2D Base Code. `http://www.rctools.sourceforge.jp` (2010)
11. Moore, E., Obst, O., Prokopenko, M., Wang, P., Held, J.: Gliders2012: Development and competition results. CoRR **abs/1211.3882** (2012)
12. Tavafi, A., Nozari, N., Vatani, R., Yousefi, M.R., Rahmatinia, S., Pirdir, P.: MarliK 2012 Soccer 2D Simulation Team Description Paper. In: RoboCup 2012 Symposium and Competitions: Team Description Papers, Mexico City, Mexico, June 2012. (2012)
13. Abdallah, M., Gama, Y.E., Antoniou, C.: robOTTO RoboCup Soccer Simulation 2D Team Description. In: RoboCup 2012 Symposium and Competitions: Team Description Papers, Mexico City, Mexico, June 2012. (2012)
14. Marian, S., Luca, D., Sarac, B., Cotarlea, O.: Oxsy 2011 team description paper. In: RoboCup 2011 Symposium and Competitions: Team Description Papers, Istanbul, Turkey, 2011. (2011)
15. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Detecting non-trivial computation in complex dynamics. In Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A., eds.: Proceedings of the 9th European Conference on Artificial Life (ECAL 2007), Lisbon, Portugal. Volume 4648 of Lecture Notes in Artificial Intelligence., Berlin / Heidelberg, Springer (2007) 895–904
16. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local information transfer as a spatiotemporal filter for complex systems. Physical Review E **77**(2) (2008) 026110
17. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Information modification and particle collisions in distributed computation. Chaos **20**(3) (2010) 037109
18. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Coherent information structure in complex computation. Theory in Biosciences **131** (2012) 193–203
19. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local measures of information storage in complex distributed computation. Information Sciences **208** (2012) 39–54
20. Schreiber, T.: Measuring information transfer. Physical Review Letters **85**(2) (2000) 461–464
21. Wang, X.R., Miller, J.M., Lizier, J.T., Prokopenko, M., Rossi, L.F.: Quantifying and tracing information cascades in swarms. PLoS ONE **7**(7) (07 2012) e40084
22. Obst, O., Boedecker, J.: Flexible coordination of multiagent team behavior using HTN planning. In Noda, I., Jacoff, A., Bredenfeld, A., Takahashi, Y., eds.: RoboCup 2005: Robot Soccer World Cup˜IX. Lecture Notes in Artificial Intelligence. Springer, Berlin, Heidelberg, New York (2006) 521–528