

Disruptive innovations in RoboCup 2D Soccer Simulation League: from Cyberoos'98 to Gliders2016

Mikhail Prokopenko¹ and Peter Wang²

¹ Complex Systems Research Group, Faculty of Engineering and IT
The University of Sydney, NSW 2006, Australia

mikhail.prokopenko@sydney.edu.au

² Data Mining, CSIRO Data61
PO Box 76, Epping, NSW 1710, Australia

Abstract. We review disruptive innovations introduced in the RoboCup 2D Soccer Simulation League over the twenty years since its inception, and trace the progress of our champion team (Gliders). We conjecture that the League has been developing as an ecosystem shaped by diverse approaches taken by participating teams, increasing in its overall complexity. A common feature is that different champion teams succeeded in finding a way to decompose the enormous search-space of possible single- and multi-agent behaviours, by automating the exploration of the problem space with various techniques which accelerated the software development efforts. These methods included interactive debugging, machine learning, automated planning, and opponent modelling. The winning approach developed by Gliders is centred on human-based evolutionary computation which optimised several components such as an action-dependent evaluation function, dynamic tactics with Voronoi diagrams, information dynamics, and bio-inspired collective behaviour.

1 Introduction

Agent Smith: "You can't win, it's pointless to keep fighting!"

Why, Mr. Anderson? Why do you persist?"

Neo: "Because I choose to."

The Matrix Revolutions.

The first official RoboCup was held in 1997, proposing a new benchmark for Artificial Intelligence (AI) and robotics. Incidentally, another classical AI challenge was successfully met in May 1997 when IBM Deep Blue defeated the human world champion in chess. By design, RoboCup and chess differ in a few key elements: environment (static vs dynamic), state change (turn-taking vs real-time), information accessibility (complete vs incomplete), sensor readings (symbolic vs non-symbolic), and control (central vs distributed) [1]. These differences are emphasised in the RoboCup 2D Soccer Simulation League [2], which quickly gained prominence, becoming one of the largest RoboCup leagues.

In this league, two teams of 12 fully autonomous software programs (called "agents") play soccer in a two-dimensional virtual soccer stadium (11 player

agents and 1 coach agent in each team), with no remote control. Each player agent receives relative and noisy input from its virtual sensors (visual, acoustic and physical) and may perform some basic actions in order to influence its environment, e.g., running, turning and kicking the ball. The coach agent receives perfect input but can communicate with the player agents only infrequently and through a fairly limited channel. The ability to simulate soccer matches without physical robots abstracts away low-level issues such as image processing and motor breakages, allowing teams to focus on the development of complex team behaviours and strategies for a larger number of autonomous agents [3, 4].

A simulated game lasts just over 10 minutes on average, and is played over a small network of computer workstations which execute the code in parallel. Each simulation step takes merely a tenth of a second, during which the entire sensory-motor cycle takes place within an agent: starting with receiving new sensory inputs from the simulator, proceeding to updating the internal memory, to evaluating possible choices, to sending the chosen action back to the simulator. The main challenge for each agent is to derive the best possible action to execute at any specific time, while facing unexpected actions of the opposing agents.

Over 20 years, the RoboCup community has developed the open-source 2D simulator and visualisation software which currently, with various packaged utilities and basic agent libraries, contains nearly a million lines of code. During this period, the League and the participating teams have undergone several transitions each of which eventually expanded the level of agents' intelligence and their behavioral complexity. In this paper we attempt to trace not only the ten-year long progress of our own team from its first implementation (Cyberroos; participated between 1998 and 2003) to the RoboCup-2016 champion team (Gliders; competed first in 2012), but also put this trace in the context of the twenty-year long evolution of the Sim2D League itself.

The conjecture we put forward is that the League has been developing as an ecosystem with an increasing complexity shaped by different approaches taken by participating teams. Furthermore, this evolving ecosystem has experienced a series of salient transitions leading to emergence of qualitatively new properties in the intelligence exhibited by the agents. By a transition we do not mean a mere extension of some simulated capabilities, such as the introduction of goalkeepers, heterogeneous player types, or a coach language. Instead, we associate a transition with a specific methodological advance which played the role of a *disruptive innovation*, with wide-spread consequences affecting the entire "ecosystem", for example, a release of standard libraries, and so on. We use the term "disruptive innovation" in a broad sense to indicate an innovation that creates a new ecosystem (by analogy with a new market or value network), eventually disrupting an existing system, displacing established structures and relationships.

2 A simulated world

The foundation supporting the evolution of the League is undoubtedly the construction of the soccer server itself, providing a centralised world model with several key features, enhanced over the following years:

- distributed client/server system running on a network and producing fragmented, localised and imprecise (noisy and latent) information about the environment (virtual soccer field) [5, 6];
- concurrent communication with a number of autonomous agents [7];
- heterogeneous sensory data (visual, auditory, kinetic) without a global vision, and limited range of basic commands/actuators (turn, kick, dash, ...) [8];
- asynchronous perception-action activity and limited window of opportunity to perform an action [9];
- autonomous decision-making under constraints enforced by teamwork (collaboration) and opponent (competition) [10];
- conflicts between reactivity and deliberation [11].

The only restriction that was imposed from the outset is that participants should “never use central control mechanisms to control a team of agents” [12].

A crucial feature making this simulated world an evolving “ecosystem” is the availability of binaries (and sometimes the source code) of participating teams, contained within an online team repository. The repository is updated after each annual RoboCup competition, allowing the participants to improve their teams with respect to the top teams of the previous championships. These improvements diversify the teams’ functionality and explore the immense search-space of possible behaviours in the quest for optimal solutions. This process results in a co-evolution of the teams, raising the overall competition level.

3 Partial automation of development efforts

“AT Humboldt” from Humboldt University, Germany became the first champion of the League at RoboCup-1997 (Nagoya, Japan). The team used a combination of reactive and planning systems, successfully deploying its agents within the simulated world.

The following couple of years passed under the domination of “CMUnited” team from Carnegie Mellon University (USA) which took the championship in 1998 (Paris, France) and 1999 (Stockholm, Sweden). One of the key reasons for this success was the development of several tools partially automating the overall effort, such as an offline agent training module, and layered disclosure: a technique for disclosing to a human designer the specific detailed reasons for an agent actions (in run-time or retroactively). Layered disclosure made it possible to inspect the details of an individual player’s decision-making process at any point [13], becoming in our view the first disruptive innovation in the League. Together with the offline agent training module, it clearly exemplified the power of automation in accelerating the development effort — precisely because it enabled the design effort to reach into a larger part of the search-space by encoding more diverse behaviours.

It is important to point out that there were other novelties introduced by CMUnited-98 and CMUnited-99, such as “single-channel, low-bandwidth communication”, “predictive, locally optimal skills (PLOS)”, “strategic positioning using attraction and repulsion (SPAR)”, etc. [13], but we believe that it is the partial automation of the software development that became the *disruptive* innovation. It has led to a wide-spread adoption of several debugging, visualising, log-playing, log-analysing, and machine learning tools.

4 Configurational space

A number of new teams in 2000 utilised the code base of the 1999 champions, CMUnited-99: it provided code for interaction with the soccer server, skills, strategies, and debugging tools in a variety of programming languages [14]. The champion of RoboCup-2000 held in Melbourne, Australia, “FC Portugal” from University of Aveiro and University of Porto, extended this code base with a systematic approach to describing team strategy, the concepts of tactics, formations and player types, as well as the situation based strategic positioning, the dynamic positioning and role exchange mechanisms [15, 11].

The generic innovation underlying these mechanisms comprised the ability to configure diverse single- and multi-agent behaviours. The range of these behaviours span from active (ball possession) to strategic (ball recovery), from formations to tactics, and from individual skills to team strategies. Such diversity resulted in a considerable configurational flexibility displayed by the winning team, significantly increasing the software development productivity, and more importantly, expanding the extent of the available behavioural search-space.

Not surprisingly, the expansion brought about by the larger configurational capacity was further exploited by the introduction of a standard coach language [16] enabling high-level coaching with explicit definition of formations, situations, player types and time periods, and resulting in a high-level coordination of team behaviour. In other words, a disruptive innovation again was delivered by a method which allowed to access deeper regions of the available search-space.

Team “TsinghuAeolus” from Tsinghua University, China, which won the next two championships (RoboCup-2001 in Seattle, USA, and RoboCup-2002 in Fukuoka), focussed specifically on increasing the agents’ adaptability via a novel online advice-taking mechanism [17]. The configurational space was extended by a task-decomposition mechanism that assigned different parts of the task to different agents.

A major boost to the League was provided by the partial release of the source code of the next champion, team “UvA Trilearn” from University of Amsterdam, The Netherlands, which won RoboCup-2003 in Padua, Italy [18]. This release resulted in a standardisation of many low-level behaviours and world model, effectively “locking in” the configurational space attained by that time, and motivating several teams to switch their code base to UvA Trilearn base.

5 Cyberoos: 1998 – 2003

At this stage we take a brief look at our first team, Cyberoos, which participated in RoboCup competitions between 1998 and 2003. The Cyberoos’98 team took 3rd place in the 1998 Pacific Rim RoboCup competition [19], while Cyberoos’2000 were 4th in the Open European RoboCup-2000 [9]. Despite these regional successes, the team’s best result at the world stage was a shared 9th place which Cyberoos repeatedly took at the RoboCup competitions in 2000, 2001, 2002 and 2003, never reaching the quarter-finals [20–23]. In hindsight, the main reason for this lack of progress was an oversight of the main tendency driving the innovations in the League: the exploration of the search-space due

to the automation of the development efforts and the standardisation of the configurational space.

Instead, the approach taken by Cyberoos focussed on self-organisation of emergent behaviour within a purely reactive agent architecture [21]. Only during the later years the Cyberoos architecture diversified, and included semi-automated methods that quantified the team performance in generic information-theoretic terms [22, 23]. This approach focussed on measuring the behavioural and belief dynamics in multi-agent systems, offering a possibility to evolve the team behaviour, optimised under a universal objective function, within the framework of information-driven self-organisation [24–26]. However, this framework has started to take a functional shape only a few years later, after the time when the Cyberoos team effort stopped in 2003.

6 Search-space decomposition

The next decade of RoboCup championships witnessed an intense competition between three teams: “Brainstormers” from University of Osnabrück, Germany, “WrightEagle” from University of Science and Technology of China, and “HELIOS” from Fukuoka University and Osaka Prefecture University, Japan. Brainstormers became champions three times: in 2005 (Osaka, Japan), 2007 (Atlanta, USA), and 2008 (Suzhou, China); WrightEagle came first an incredible six times: in 2006 (Bremen, Germany), 2009 (Graz, Austria), 2011 (Istanbul, Turkey), 2013 (Eindhoven, The Netherlands), 2014 (Joao Pessoa, Brazil) and 2015 (Hefei, China); and HELIOS succeeded twice: in 2010 (Singapore) and 2012 (Mexico City, Mexico).

6.1 Machine learning

Brainstormers’ effort focussed on reinforcement learning methods aiming at a universal machine learning system, where the agents learn to generate the appropriate behaviors to satisfy the most general objective of “winning the match”. Unfortunately, as has been acknowledged [27], “even from very optimistic complexity estimations it becomes obvious, that in the soccer simulation domain, both conventional solution methods and also advanced today’s reinforcement learning techniques come to their limit – there are more than $(108 \times 50)^{23}$ different states and more than $(1000)^{300}$ different policies per agent per half time”.

The high dimensionality of the search space motivated Brainstormers to use a multilayer perceptron neural network [27]: a feedforward artificial neural network which utilises a supervised learning technique called backpropagation for training the network. Rather than developing a universal learning system, Brainstormers succeeded in decomposing the problem into a number of individual behaviours (e.g., NeuroKick, NeuroIntercept, NeuroHassle) and tactics (e.g., NeuroAttack2vs2, NeuroAttack3vs4, NeuroAttack7vs8), learned with supervised learning techniques.

Recently, there has been some renewed interest in backpropagation networks due to the successes of deep learning. In our view, the potential of reinforcement learning methods in RoboCup has not yet been fully realised, and deep learning may yet to become a disruptive innovation for the Simulation league.

6.2 Automated planning

WrightEagle team addressed the challenges of (i) high dimensionality of the search space and (ii) the limited computation time available in each decision cycle, by using Markov Decision Processes (MDPs). The developed framework decomposes a given MDP into a set of sub-MDPs arranged over a hierarchical structure, and includes heuristics approximating online planning techniques [28]. WrightEagle approach abandoned “the pursuit of absolute accuracy” and divided the continuous soccer field into the discrete space, further subdividing it into the players’ control areas according to geometric reachability. The resultant structure enables automated planning, accelerating the search process and extending the search depth [28].

6.3 Opponent modelling

“HELIOS” team [29, 30] followed a similar path, targeting a decomposition of the problem space in developing an unsupervised learning method based on Constrained Delaunay Triangulation (CDT) [31]. A Delaunay triangulation for a set P of points in a plane is a triangulation $\mathcal{D}(P)$ such that no point in P is inside the circumcircle of any triangle in $\mathcal{D}(P)$ (in CDT the circumcircle of some triangles contains other triangles’ vertices). The method divides the soccer field into a set of triangles, which provide an input plane region for Neural Gas (NG) and Growing Neural Gas (GNG) methods. Specifically, the set P_b of N points represents specifically chosen positions of the ball on the field, while sets P_i describe the sets of coordinates of each player $1 \leq i \leq 11$, so that there is a bijective correspondence between P_b and each of P_i . Moreover, when the ball takes any position within a triangle of $\mathcal{D}(P_b)$, each player’s position is computed in a congruent way within $\mathcal{D}(P_i)$. During offline experiments or even during a game, the behaviour of the opponent, for example, the players’ motion, directions of the passes, and the overall team formations, can be mapped, analysed and categorised [29, 30].

It is evident that the main reason behind the recurrent successes of all three champion approaches is a dynamic decomposition of the problem space and its subsequent efficient exploration. This innovation goes beyond a simple standardisation of low-level behaviours within a rich but static configurational space, by employing automated learning and planning methods in a dynamic search.

7 Standardisation of “hardware”

An influential disruptive innovation arrived in 2010, when HELIOS team released a major update of their well-developed code base [32]:

- *librcsc-4.0.0*: a base library for the RoboCup Soccer Simulator (RCSS);
- *agent2d-3.0.0*: a base source code for a team;
- *soccerwindow2-5.0.0*: a viewer and a visual debugger program for RCSS;
- *fedit2-2.0.0*: a team formation editor for *agent2d*.

This resulted in nearly 80% of the League’s teams switching their code base to *agent2d* over the next few years. One may think of this phenomenon as a standardisation of the simulated hardware, freeing the effort to improving the higher-level tactical behaviours.

8 Gliders (2012 – 2016): Fusing human innovation and artificial evolution

We turn our attention to our champion team which won RoboCup-2016 (Leipzig, Germany): Gliders [33–37]. Gliders2012 and Gliders2013 reached the semi-finals of RoboCup in 2012 and 2013; Gliders2014 were runner-ups in 2014; Gliders2015 finished third in RoboCup-2015, and Gliders2016 (a joint effort of the University of Sydney and CSIRO) became world champions in 2016.

RoboCup-2016 competition included 18 teams from 9 countries: Australia, Brazil, China, Egypt, Germany, Iran, Japan, Portugal and Romania. Gliders2016 played 23 games during several rounds, winning 19 times, losing twice and drawing twice, with the total score of 62:13, or 2.70 : 0.57 on average. In the two-game semi-final round, Gliders2016 defeated team CSU_Yunlu from Central South University (China), winning both games with the same score 2:1. The single-game final against team HELIOS2016 (Japan) went into the extra time, and ended with Gliders2016 winning 2:1. The third place was taken by team Ri-one from Ritsumeikan University (Japan).

The 2016 competition also included an evaluation round, where all 18 participating teams played one game each against the champion of RoboCup-2015, team WrightEagle (China). Only two teams, the eventual finalists Gliders2016 and HELIOS2016, managed to win against the previous year champion, with Gliders defeating WrightEagle 1:0, and HELIOS producing the top score 2:1.

The Gliders team code is written in C++ using *agent2d-3.1.1* [32], and fragments of source code of team MarliK released in 2012 [38].

In order to optimise the code, the Gliders development effort over the last five years involved *human-based evolutionary computation* (HBEC): a set of evolutionary computation techniques that rely on human innovation [39, 40].

In general, evolutionary algorithms search a large space of possible solutions that together form a population. Each solution is a “genotype”: a complex data structure representing the entire team behaviour encoded through a set of “design points”. A design point can be as simple as a single parameter (e.g., risk tolerance in making a pass), or as complicated as a multi-agent tactical behaviour (e.g., a conditional statement describing the situation when a defender moves forward to produce an offside trap).

Some design points are easy to vary. For instance, a formation defined via Delaunay Triangulations $\mathcal{D}(P_b)$ and $\mathcal{D}(P_i)$, $1 \leq i \leq 11$, is an ordered list of coordinates, and varying and recombining such a list can be relatively easily automated. Other design points have an internal structure and are harder to permute. For example, a conditional statement describing a tactic has a condition and an action, encoded by numerous parameters such as positional coordinates, state information, and action details. Once such a statement (a design point) is created by human designers, its encoding can be used by evolutionary algorithms. However, the inception of the tactic needs creative innovation in the first place, justifying the hybrid HBEC approach.

The HBEC solutions representing team behaviours are evaluated with respect to their fitness, implemented as the average team performance, estimated over thousands of games for each generation played against a specific opponent.

Some solutions are retained and recombined (i.e., the members of the population live) and some are removed (i.e., die) through selection. Importantly, the evolutionary process is carried out within different landscapes (one per known opponent), and typically results in different solutions evolved to outperform specific opponents. In order to maintain coherence of the resultant code, each design point is implemented with a logical mask switching the corresponding part of the genotype on and off for specific opponents (determined by their team names). This is loosely analogous to epigenetic programming [41].

The approach is aimed at constantly improving performance from one artificial “generation” to another, with team designers innovating and recombining behaviours while the fitness landscape and the mutations are for the most part automated. The performance of Gliders was evaluated on several supercomputer clusters, executing on some days tens of thousands of the experimental runs with different behaviour versions. It would be a fair estimate that the number of such trials is approaching 10 million. The overall search-space explored by the HBEC includes variations in both Gliders behaviour and opponent modelling. The approach incorporates disruptive innovations of the past years, including the standardisation of simulated “hardware” and several effective search-space decompositions.

Specific variations included (i) action-dependent evaluation function, (ii) dynamic tactics with Voronoi diagrams, (iii) information dynamics, and (iv) bio-inspired collective behaviour.

The approach introduced in Gliders2012 [33] retained the advantages of a single evaluation metric (implemented in *agent2d* [32]), but diversified the evaluation by considering multiple points as desirable states. These desirable states for action-dependent evaluation are computed using Voronoi diagrams which underlie many tactical schemes of Gliders.

Starting from 2013, Gliders utilised information dynamics [42–47] for tactical analysis and opponent modelling. This analysis involves computation of information transfer and storage, relating the information transfer to responsiveness of the players, and the information storage within a team to the team’s rigidity and lack of tactical richness.

The constraints on mobility, identified by the information dynamics, were investigated and partially overcome with bio-inspired collective behaviour [36]. Gliders2015 utilise several elements of swarm behavior, attempting to keep each player’s position as close as possible to that suggested by a specific tactical scheme, while incorporating slight variations in order to maximise the chances of receiving the pass and/or shooting at the opponent’s goal. This behaviour increased the degree of coherent mobility: on the one hand, the players are constantly refining their positions in response to opponent players, but on the other hand, the repositioning is not erratic and the players move in coordinated ways.

These directions were unified within a single development and evaluation framework which allowed to explore the search-space in two ways: translating human expertise into new behaviours and tactics, and exhaustively recombining them with an artificial evolution, leveraging the power of modern supercomputing. This fusion, we believe, produced a disruptive innovation on its own, providing the winning edge for Gliders.

9 Conclusion

In this paper we reviewed disruptive innovations which affected advancement of the RoboCup 2D Soccer Simulation League over the twenty years since its inception, and placed the progress of our champion team in this context. It is important to realise that the neither of these processes has been linear, and many ideas have been developing along a spiral-shaped trajectory, resurfacing over the years in a different implementation. For example, the utility of evolutionary computation supported by supercomputing has been suggested as early as 1997, when a simulated team was developed with the agents whose high-level decision making behaviors had been entirely evolved using genetic programming [48]. Yet the complexity of the domain proved to be too challenging for this approach to gain a widespread adoption at that time.

Without an exception, all the winning approaches combined elements of some automation (debugging, machine learning, planning, opponent modelling, evolutionary computation) with human-based innovation in terms of a decomposition of the search-space, providing various configurations, templates and structures. Is there still a way toward a fully automated solution, when the agents learn or evolve to play a competitive game without a detailed guidance from human designers, but rather by trying to satisfy a universal objective (“win a game”)?

On the one hand, the ability to run a massive number of simulated games on supercomputing clusters producing replicable results will only strengthen in time [4], and so may lend some hope in meeting this challenge positively. On the other hand, the enormous size and dimensionality of the search-space would defy any unstructured exploration strategy. A methodology successfully resolving this dilemma may not only provide an ultimate disruptive innovation in the League, but also provide a major breakthrough in the general AI research.

Acknowledgments

Several people contributed to Cyberoos and Gliders development over the years. Marc Butler, Thomas Howard and Ryszard Kowalczyk made exceptionally valuable contributions to Cyberoos’ effort during 1998 – 2002 [19, 9, 20, 21]. We are grateful to Gliders team members Oliver Obst, particularly for establishing the tournament infrastructure supporting the team’s performance evaluation on CSIRO Accelerator Cluster (Bragg), and Victor Jauregui, for several important insights on soccer tactics used in Gliders2016 [37]. We thank David Budden for developing a new self-localisation method introduced in Gliders2013 [49, 34] as well as contributing to the analysis of competition formats [4], and Oliver Cliff for developing a new communication scheme adopted by Gliders from 2014 [35]. The overall effort has also benefited from the study quantifying tactical interaction networks, carried out in collaboration with Oliver Cliff, Joseph T. Lizier, X. Rosalind Wang and Oliver Obst [46]. We are thankful to Ivan Duong, Edward Moore and Jason Held for their contribution to Gliders2012 [33]. Gliders team logo was created by Matthew Chadwick.

References

1. Asada, M., Kitano, H., Noda, I., Veloso, M.: RoboCup: Today and tomorrow – What we have have learned. *Artificial Intelligence* **110** (1999) 193–214

2. Kitano, H., Tambe, M., Stone, P., Veloso, M.M., Coradeschi, S., Osawa, E., Matsubara, H., Noda, I., Asada, M.: The RoboCup Synthetic Agent Challenge 97. In: *RoboCup-97: Robot Soccer World Cup I*. Springer, London (1998) 62–73
3. Budden, D., Wang, P., Obst, O., Prokopenko, M.: Simulation leagues: Analysis of competition formats. In Bianchi, R.A.C., Akin, H.L., Ramamoorthy, S., Sugiura, K., eds.: *RoboCup 2014: Robot Soccer World Cup XVIII*. Volume 8992 of *Lecture Notes in Computer Science.*, Springer (2014) 183–194
4. Budden, D.M., Wang, P., Obst, O., Prokopenko, M.: Robocup simulation leagues: Enabling replicable and robust investigation of complex robotic systems. *IEEE Robotics and Automation Magazine* **22**(3) (2015) 140–146
5. 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
6. Haker, M., Meyer, A., Polani, D., Martinetz, T.: A method for incorporation of new evidence to improve world state estimation. In Birk, A., Coradeschi, S., Tadokoro, S., eds.: *RoboCup 2001: Robot Soccer World Cup V*. Volume 2377 of *Lecture Notes in Computer Science.*, Springer (2002) 362–367
7. Stone, P., Veloso, M.: Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence* **110**(2) (1999) 241–273
8. 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. *Lecture Notes in Artificial Intelligence*. Springer (2001)
9. Butler, M., Prokopenko, M., Howard, T.: Flexible synchronisation within RoboCup environment: A comparative analysis. In Stone, P., Balch, T., Kraetzschmar, G., eds.: *RoboCup 2000: Robot Soccer World Cup IV*. Volume 2019 of *Lecture Notes in Artificial Intelligence*. Springer, Berlin, Heidelberg (2001) 119–128
10. 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)
11. 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*. Springer, London, UK (2001) 175–197
12. Noda, I., Suzuki, S., Matsubara, H., Asada, M., Kitano, H.: Robocup-97: The first robot world cup soccer games and conferences. *AI Magazine* **19**(3) (1998) 49–59
13. Stone, P., Riley, P., Veloso, M.: The CMUnited-99 champion simulator team. In Veloso, M., Pagello, E., Kitano, H., eds.: *RoboCup-99: Robot Soccer World Cup III*. Volume 1856 of *Lecture Notes in Artificial Intelligence*. Springer (2000) 35–48
14. Stone, P., Asada, M., Balch, T.R., Fujita, M., Kraetzschmar, G.K., Lund, H.H., Scerri, P., Tadokoro, S., Wyeth, G.: Overview of RoboCup-2000. In Stone, P., Balch, T.R., Kraetzschmar, G.K., eds.: *RoboCup 2000: Robot Soccer World Cup IV*. Volume 2019 of *Lecture Notes in Computer Science.*, Springer (2000) 1–28
15. Reis, L.P., Lau, N.: FC Portugal Team Description: RoboCup 2000 Simulation League Champion. In Stone, P., Balch, T.R., Kraetzschmar, G.K., eds.: *RoboCup 2000: Robot Soccer World Cup IV*. Volume 2019 of *Lecture Notes in Computer Science.*, Springer (2001) 29–40
16. Reis, L., Lau, N.: COACH UNILANG - A Standard Language for Coaching a (Robo)Soccer Team. In Birk, A., Coradeschi, S., Tadokoro, S., eds.: *Robocup 2001: Robot Soccer World Cup V*. Volume 2377 of *Lecture Notes in Computer Science.*, Springer (2001) 183–192

17. Jinyi, Y., Ni, L., Fan, Y., Yunpeng, C., Zengqi, S.: Technical Solutions of TsinghuAeolus for Robotic Soccer. In Polani, D., Browning, B., Bonarini, A., Yoshida, K., eds.: *RoboCup 2003: Robot Soccer World Cup VII*, Springer (2004) 205–213
18. Kok, J.R., Vlassis, N., Groen, F.: UvA Trilearn 2003 team description. In Polani, D., Browning, B., Bonarini, A., Yoshida, K., eds.: *Proceedings CD RoboCup 2003*, Padua, Italy, Springer (July 2003)
19. Prokopenko, M., Kowalczyk, R., Lee, M., Wong, W.Y.: Designing and modelling situated agents systematically: Cyberoos98. In: *Proceedings of the PRICAI-98 Workshop on RoboCup*. (1998) 75–89
20. Prokopenko, M., Butler, M., Howard, T.: On emergence of scalable tactical and strategic behaviour. In Stone, P., Balch, T., Kraetzschmar, G., eds.: *RoboCup 2000: Robot Soccer World Cup IV*. Volume 2019 of *Lecture Notes in Artificial Intelligence*. Springer, Berlin, Heidelberg (2001) 357–366
21. Prokopenko, M., Wang, P., Howard, T.: Cyberoos'2001: "Deep Behaviour Projection" Agent Architecture. In Birk, A., Coradeschi, S., Tadokoro, S., eds.: *RoboCup 2001: Robot Soccer World Cup V*. Volume 2377 of *Lecture Notes in Computer Science*., Springer (2002) 507–510
22. 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
23. 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 (2004) 89–101
24. Nehaniv, C., Polani, D., Olsson, L., Klyubin, A.: Evolutionary information-theoretic foundations of sensory ecology: Channels of organism-specific meaningful information. *Modeling Biology: Structures, Behaviour, Evolution* (2005) 9–11
25. Prokopenko, M., Gerasimov, V., Tanev, I.: Measuring spatiotemporal coordination in a modular robotic system. In Rocha, L., Yaeger, L., Bedau, M., Floreano, D., Goldstone, R., Vespignani, A., eds.: *Artificial Life X: Proceedings of The 10th International Conference on the Simulation and Synthesis of Living Systems*, Bloomington IN, USA (2006) 185–191
26. Prokopenko, M., Gerasimov, V., Tanev, I.: Evolving spatiotemporal coordination in a modular robotic system. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J.C.T., Marocco, D., Meyer, J.A., Miglino, O., Parisi, D., eds.: *From Animals to Animats 9: 9th International Conference on the Simulation of Adaptive Behavior (SAB 2006)*, Rome, Italy, September 25-29 2006. Volume 4095 of *Lecture notes in computer science*. (2006) 558–569
27. Riedmiller, M., Gabel, T., Trost, F., Schwegmann, T.: Brainstormers 2D – Team Description 2008. In: *RoboCup 2008: Robot Soccer World Cup XII; CD*. (2008)
28. Zhang, H., Chen, X.: The decision-making framework of WrightEagle, the RoboCup 2013 soccer simulation 2D league champion team. In: *Robot Soccer World Cup*, Springer (2013) 114–124
29. Akiyama, H., Noda, I.: Multi-agent positioning mechanism in the dynamic environment. In Visser, U., Ribeiro, F., Ohashi, T., Dellaert, F., eds.: *RoboCup 2007: Robot Soccer World Cup XI*. Springer, Berlin, Heidelberg (2008) 377–384
30. Akiyama, H., Shimora, H.: Helios2010 team description. In: *RoboCup 2010: Robot Soccer World Cup XIV; CD*. (2010)
31. Chew, L.P.: Constrained Delaunay Triangulations. *Algorithmica* 4(1-4) (1989) 97–108
32. Akiyama, H.: Agent2D Base Code. <http://www.rctools.sourceforge.jp> (2010)

33. 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)
34. Prokopenko, M., Obst, O., Wang, P., Budden, D., Cliff, O.: Gliders2013: Tactical analysis with information dynamics. In: RoboCup 2013 Symposium and Competitions: Team Description Papers, Eindhoven, The Netherlands, June 2013. (2013)
35. Prokopenko, M., Obst, O., Wang, P.: Gliders2014: Dynamic Tactics with Voronoi Diagrams. In: RoboCup 2014 Symposium and Competitions: Team Description Papers, Joao Pessoa, Brazil, July 2014. (2014)
36. Prokopenko, M., Wang, P., Obst, O.: Gliders2015: Opponent avoidance with bio-inspired flocking behaviour. In: RoboCup 2015 Symposium and Competitions: Team Description Papers, Hefei, China, July 2015. (2015)
37. Prokopenko, M., Wang, P., Obst, O., Jaurgeui, V.: Gliders2016: Integrating multi-agent approaches to tactical diversity. In: RoboCup 2016 Symposium and Competitions: Team Description Papers, Leipzig, Germany, July 2016. (2016)
38. 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. (2012)
39. Kosorukoff, A.: Human based genetic algorithm. In: Systems, Man, and Cybernetics, 2001 IEEE International Conference on. Volume 5., IEEE (2001) 3464–3469
40. Cheng, C.D., Kosorukoff, A.: Interactive one-max problem allows to compare the performance of interactive and human-based genetic algorithms. In Deb, K., ed.: Genetic and Evolutionary Computation – GECCO 2004: Genetic and Evolutionary Computation Conference, Seattle, USA, June 26–30, 2004. Springer (2004) 983–993
41. Tanev, I., Yuta, K.: Epigenetic programming: Genetic programming incorporating epigenetic learning through modification of histones. *Information Sciences* **178**(23) (2008) 4469–4481
42. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Information modification and particle collisions in distributed computation. *Chaos* **20**(3) (2010) 037109
43. 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) (2012) e40084
44. Ay, N., Bernigau, H., Der, R., Prokopenko, M.: Information-driven self-organization: the dynamical system approach to autonomous robot behavior. *Theory in Biosciences* **131** (2012) 161–179
45. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Coherent information structure in complex computation. *Theory in Biosciences* **131** (2012) 193–203
46. Cliff, O., Lizier, J., Wang, R., Wang, P., Obst, O., Prokopenko, M.: Towards quantifying interaction networks in a football match. In Behnke, S., Veloso, M., Visser, A., Xiong, R., eds.: RoboCup 2013: Robot Soccer World Cup XVII, Springer (2013) 1–12
47. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: A framework for the local information dynamics of distributed computation in complex systems. In Prokopenko, M., ed.: Guided Self-Organization: Inception. Springer (2014) 115–158
48. Luke, S.: Genetic Programming Produced Competitive Soccer Softbot Teams for RoboCup97. In Koza, J.R., Banzhaf, W., Chellapilla, K., Deb, K., Dorigo, M., Fogel, D.B., Garzon, M.H., Goldberg, D.E., Iba, H., Riolo, R.L., eds.: Proceedings of the 3rd Annual Genetic Programming Conf., Morgan Kaufmann (1998) 214–222
49. Budden, D., Prokopenko, M.: Improved particle filtering for pseudo-uniform belief distributions in robot localisation. In Behnke, S., Veloso, M.M., Visser, A., Xiong, R., eds.: RoboCup 2013: Robot Soccer World Cup XVII. Volume 8371 of Lecture Notes in Artificial Intelligence., Springer (2013)