Coordination and Cooperation in Robot Soccer

Vincenzo Suriani[1] §[0000-0003-1199-8358],
Emanuele Antonioni[1] §[0000-0001-8863-9076],
Francesco Riccio[2][0000-0002-9112-8143], and Daniele Nardi[1][0000-0001-6606-200X]

Department of Computer, Control, and Management Engineering
Sapienza University of Rome, Rome (Italy)
{lastname}@diag.uniroma1.it.

Abstract. Aiming at improving our physical strength and expanding our knowledge, tournaments and competitions have always contributed to our personal growth. Robotics and AI are no exception, and since beginning, competitions have been exploited to improve our understanding of such research areas (e.g. Chess, VideoGames, DARPA). In fact, the research community has launched (and it is involved) in several robotics competitions that provide a two-fold benefit of (i) promoting novel approaches and (ii) valuate proposed solutions systematically and quantitatively. In this paper, we focus on a particular research area of Robotics and AI: we analyze multi-robot systems deployed in a cooperative-adversarial environment being tasked to collaborate to achieve a common goal, while competing against an opposing team. To this end, RoboCup provide the best benchmarking environment by implementing such a challenging problem in the game of soccer. Sports, in fact, represent extremely complex challenge that require a team of robots to show dexterous and fluid movements and to feature high-level cognitive capabilities. Here, we analyse methodologies and approaches to address the problem of coordination and cooperation and we discuss state-of-the-art solutions that achieve effective decision-making processes for multi-robot adversarial scenarios.

Keywords: Strategies in Robotic Games · Robotic Competition · Soccer Robots RoboCup SPL

1 Introduction

Games and sport competitions offer a suitable application where both teammates cooperation and opponents management play a key role. Hence, being able to deploy a artificial agent capable of showing human (or even super-human) performance in these contexts, is one of the most difficult and fascinating goal that lies at the intersection of robotics and artificial intelligence. Several milestones have already be reached in this race for progress and technological advancement. One of the first steps, most known to the mass, is undoubtedly the chess playing AI that beat the human world-champion for the first time, DeepBlue. More

§ These two authors contributed equally.

recently, but still in the context of board-games, we acknowledge AlphaGo [23] that, similarly to its predecessor, beat the world-champion in the most complex board-game in history, the game of GO. Moreover, the techniques investigated in recent years show promising results being generalizable to different scenarios and agnostic to the state representation. In fact, more related to the work we address in this survey, we also want to report the breakthrough achieved in [17] where the authors successfully beat a team of humans in an highly-interactive, partially observable, multi-agent scenarios in a continuous state-space world.

Assuming a different perspective in the wide spectrum of multi-agent approaches, here, we investigate how the research community tackles the problem of deploying such techniques on real robots playing soccer. In particular, we explore: (i) individual strategies in multi-agent scenario; (ii) cooperative and strategic decision-making; and (iii) opponent behavior analysis in adversarial settings. Operating in the real world adds new challenges and complexity to problems to solve. In fact, in this context, proposed techniques have to necessarily take into account system failures, noisy perceptions, unpredictable and non-stationary environments, and numerous unknown events ranging from faulty physical components to opponents high-level strategies. Hence, in order to highlight the most promising techniques and determine next research directions, we believe that categorizing the main contributions implemented within the RoboCup competitions is key to provide a solid basis in the deployment of state-of-the-art approaches to coordination and cooperation on physical robot.

It is very difficult to develop and deploy an autonomous agent able to understand and act in the physical world. In fact, when operating in uncontrolled scenarios, robots must show robust and effective skills to support perception, reasoning and coordinated behaviors with people and of course, other agents. To this end, the research community is constantly promoting robotic competitions in order to solve particular tasks and to develop and deploy operating agents in the real world. In this context, RoboCup is one of the leadership organization that challenges participants world-side in the game of soccer [4] with the aim to develop and end-to-end robotic system capable of perceiving the environment, high-level reasoning and performing agile and smooth motions.

In this context, perception and reasoning are enabling factors to enter the soccer field, but coordinated effective robot behaviors are the key factor for winning.

Each RoboCup league is designed to address a particular challenge in developing and deploying a fully autonomous robot soccer player (see Fig. 1). In fact, tackling sports at once is extremely difficult and attempting to tackle all the research questions at the same time leads to unpractical and under-performing systems. Hence, each of the RoboCup leagues is carefully defined to operate in a particular research area – even though a certain amount of overlap is guaranteed. Such an organization allows to divide the soccer game in sub-problems and to better formalize solutions for each of them. Usually, we can categorize proposed approaches in accordance with the sense-plan-act paradigm, and intuitively, each league mainly targets one of these macro areas. In this paper, we categorize the

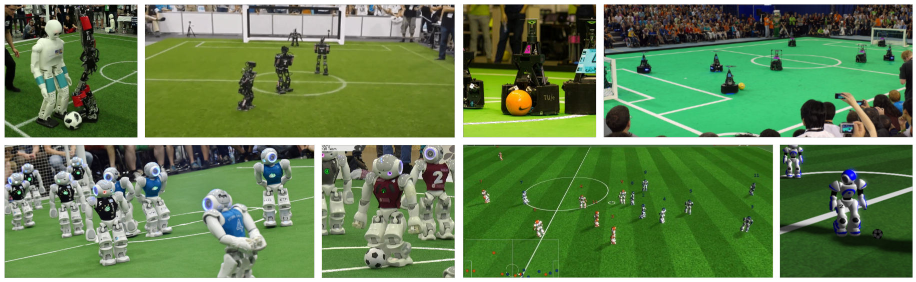

Fig. 1. RoboCup Soccer Leagues.

contributions made in the context of RoboCup with a particular focus on coordination and collective behavior across different leagues. Our goal is to understand the most competitive methodologies currently used, and to highlight the most promising trends of research that will guide us to implement a fully autonomous team of robots. Our focus is to survey proposed solutions that contribute in enabling a team of robots to collectively perceive the world, asynchronously reason on current state of the environment and optimally coordinate their action to achieve a common goal while competing with other robots.

2 RoboCup Leagues and Organization

RoboCup competitions are organized in several leagues each of which aims at tackling a particular research challenge. In this paper we focus on coordination and cooperation approaches which are a characterizing aspect of RoboCup soccer leagues. Our goal is to analyze proposed techniques in order to highlight research trends and understand their enabling factor. Such leagues are particularly suitable to advance in our understanding of multi-agent systems. In fact, the sport of soccer forces the team of robots to demonstrate robust individual and collective behaviors while competing against another teams in an adversarial setting.

However, solving the game of soccer at once is not an easy task and the organization split the problem in different research areas, each of which is assigned to a particular league. Hence, such leagues features their own challenges being designed with different environmental and structural assumptions. One of the most sharp categorization that affects the methodologies proposed in the competitions is determined by the physical implementation of the agents, i.e. the platform hardware. Leagues, in fact, range from simulated agents to heavy-hardware platforms. *Simulation2D* (Sim2D) and *Simulation3D* (Sim3D) are the less hardware-demanding leagues which makes them the most suitable scenario to promote research in designing complex collective behavior at scale. The *Small Size* (SSL) and the *Middle Size* (MSL) represent the first gate to physical agents. These leagues employ wheeled robots which alleviate locomotion constraints and are capable of performing dexterous maneuvers at high speed. Then, the *Standard Platform League* (SPL) forces all participants to use the same robotic platform, that currently is the Aldebaran humanoid NAO robot. Such a setting allows researchers to focus more on the behavior of the different agent rather than their hardware components. However, it includes in the challenge noisy perception, partial observability and bipedal locomotion. Finally, *Humanoid* Leagues (HL) represent the most hardware-demanding configuration. In this leagues robot can be 1.6 meters tall, teams are completely in charge of the hardware components and engineering smooth and agile movements. Intuitively, however, the decision-making and cognitive behaviors are less demanding and games features a maximum of 2 vs. 2 robots.

3 Cooperation Strategies

Robots involved in the RoboCup competitions are designed to understand the external world and to exhibit robust and effective behaviors. In soccer, to this end, an agent has to show individual decision-making skills to (i) promptly react local situations [13]; (ii) reason at the collective level with other teammates in order to efficiently achieve a common goal [8]; and (iii) acknowledge opposing agents in the environment that act against [5]. Accordingly, we structure our discussion in three subsections – each of them describing specific problems to be solve in these areas and relating exiting work.

3.1 From Individual to Collective Strategies

An effective behavior for an autonomous multi-agent system is strictly related to the single agents capabilities. In fact, one of the requirements to build an effective multi-agent behavior, is a stable single player behavior. If we take a closer look to the single-agents, we can classifying their behaviors into two set of categories based on their abstraction level: skills and behaviors. Skills execute primitive actions that are usually related to the core motions of the agent, while behaviors determine how to select those primitive actions to achieve a specific goal.

In multi-agent adversarial settings, *Individual strategies* are a key factor for achieving success. To feature competitive behaviors, a robot must reconstruct a model of the world by relying on its local perceptions. Then it has to feature a robust decision-making system to determine the next set of actions to perform in order to reach a given goal. To this end, a single robot has to be capable of performing dexterous low-level motions while executing sophisticated high-level behaviors.

A low-level skill is usually defined as a predefined command for robot actuators to implement action primitive. In the soccer context these are represented by the routines for kicking, passing, dribbling, diving and getting-up. Individual behaviors are generated by composing skills and/or recursively including individual behaviors [31]. Behavior design, however, has to take into account different aspects characterizing the physical robot platforms, environmental constraints and task specifics. In RoboCup, researchers investigate a large amount of approaches and technologies in order to always show more sophisticated robot capabilities. The most common approaches are based on *state machines* [21], *planners* [9] and various learning techniques, as *Evolutionary Learning* [30], *Statistical Learning*, *Deep Learning* [18, 12].

A state-machine approach, for example, can be deployed to easily model the defender behavior of the agent that have to stop the ball to avoid the goal. On the other hand, some game situations can benefit from the use of deep reinforcement learning approaches more than a model ones. For example, during penalty-kicks or in corner-kick situations. In this kind of contexts, learning-based approaches have started to be used and have been deployed even in place of the modeled approaches, as in [3, 15], where two different statistical learning methods are adopted for solving behavioral problems. In the first one, the state evaluation of a decision-making process has been carried out by means of a Learning to Rank Algorithm. In the latter, the position of the goalkeeper agent in a MSL game has been determined with a linear regression approach. Behavior modeling has been also tackled with Reinforcement Learning approaches. In fact, in [27] within the context of the Simulation 3D, an agent has been trained to score goals without previous knowledge. This result has been achieved by means of a transfer learning system instead of the classical reward shaping approach. Within the context of the SPL, in [20] the authors addressed the problem of shaping the strategy of a defender robot adapting it to the strategy of the opponents. The method used is a combination of Monte Carlo search and data aggregation (MCSDA) that allowed to adapt the discrete-action soccer policies of the defender player. Finally, for solving the static free-kick task, a classical bandit approach has been exploited in [16].

3.2 Collective Strategies

RoboCup soccer leagues forces researchers to program robots to show effective individual and local behaviors, but also to demonstrate robust teamwork and cooperative behaviors. Suggestively, developing multi-robot decision making system is a much more complex challenge due to several factors: multiple environment perception streams, distributed world representations, dynamic role-assignment and asynchronous decision-making. Moreover, in RoboCup teammates are connected via Wi-Fi which, during games, is noisy and too prone to faulty behaviors.

Hence, in order to trade-off robustness and efficiency, and to guarantee competitive collective behaviors, researchers not only rely on the current data stream but they also provide robots with a model of the environment that can be used as a surrogate representation to embed the state of the external world [19]. Among the proposed approaches to multi-robot cooperation and collaboration, we highlight two major classes: positioning approaches and role-assignment approaches.

As the name suggests, the former category of approaches has the goal of finding the best team positioning within the field. Such a positioning can be an extremely difficult task which grows exponentially with the number of players and that is subjects to numerous factors. For example, the authors in [14] assess that the contexts in which the game is currently evolving is key to re-position teammates. For instance, team formation can be adjusted depending on whether the players are in an defensive or attacking context, or equivalently depending on the current score, players can be more aggressive or more cautious if they have to manage the opponents.

Conversely, dynamic-role assignment attempts to find the optimal mapping between a set of robots and a set of active task. For example, the authors in [6, 19] use utility functions to estimate how good a robot can perform a certain task at a given time. Utility functions are particularly suitable to evaluate and coordinate teams of heterogeneous robots acting collectively. Equivalently, the authors in [24] exploit MDP to formalize individual behavioral models and determine affinity with a set of given tasks. Differently, Catacora et al. [7] use a learning-based approach to coordinate a team of robots. Their approach shows promising results and successfully coordinates two robots in particular in-game situations (e.g. penalty-kicks). However, the computational demand of such a methodology limits its application and, at the moment, they cannot run the learned policy on large teams of robots.

3.3 Opponent Analysis for Cooperation

In adversarial multi-robot environments, having an understanding of the opponents behaviors represents a remarkable advantage. In fact in such a context, if a team of robots is able to counter opponents movements, both analyzing individual behaviors and forecasting team strategies, then it can react more precisely to the situation at hand – and thus improve the team performance. In RoboCup, we report that the majority of contributions in opponent analysis comes from leagues where perception is more reliable [26] (e.g. Sim2D, Sim3D, SSL and MSL). In such leagues we notice that opponents analysis approaches can be coarsely categorized in two classes: action sequence analysis and behavior forecasting.

The former group attempts to find patterns in the action sequences of the opponents teams in order to recognize recurring strategies. For example, in [10], the authors proposed an offline opponent action analysis approaches that processes game logs in order to extrapolate action primitives. Such primitives are then coupled with in-game states and organized in an opponents behavior tree. Yasui et al. [29], instead, formalize a dissimilarity function among state-action pairs which is then used for clustering and classification. The authors extend their work to improve computational efficiency [1] and propose an clustering algorithm to analyze the agents behaviors online and promptly react by position in order to prevent the opponents to score [2].

The latter group aims at solving a forecasting task. In other words, given the current state of teammates and sequence of opponent actions, the goal is to predict the intentions of other agents and the next state of the environment.

Such a capability is key to anticipate opponents' intentions and gain a substantial advantage on them. In this setting, Li et al. [11] proposes a fuzzy inference system to classify particular in-game situation (e.g. corner-kick, passing); predict opponent trajectories; and re-position the team formation accordingly. They show that by inferring opponents intentions it is possible to double the number of won games. Similarly, the authors in [25] introduce FOSSE, a deep model-free approach that given state representation attempts to learn a transition model, forecast future states of the environment and (as in the previous case) adjusts the team formation. Finally, in [22] the authors achieve an important milestone by rolling out a learning approach on a humanoid robot in the SPL. The authors introduce SAFEL, a real-time learning-based algorithm that is capable of generating an opponent behavioral model of an agent, and counter-react strategically.

4 Analysis and Classification of the Proposed Approaches

In this paper, we survey the implementation of the different approaches to decision-making, both at the individual and collective level. In particular, we analyzed how single agent skills, coordinated team actions and opponent analysis can contribute to implement effective multi-robot systems in cooperative-adversarial scenario characterized by different specific challenges (e.g. used platform, low communication, scalable behaviors).

However, there are different considerations that can be done in order to better classify the proposed approaches in this particularly challenging context. Such categorizations provide a thorough comparison and slice the state-of-the-art along different perspectives:

Centralized vs. Distributed. Existing solutions to multi-robot coordination include a staggering amount of different techniques, each of which comes with its own advantages and disadvantage. Most contributions to the field of multi-robot coordination depend on the environmental configuration. For example, in low-bandwidth and noisy communication scenarios, a fully distributed approach is typically to be preferred in order to allow robots to act individually even though the information about teammates is outdated – or simply not coming in. Conversely, if the overall setup is characterized by a reliable communication, a centralized approach is implemented to guarantee robustness and optimally coordinated robots (e.g. Kiva system [28]).

In RoboCup, the technical committee enforces noisy and non-constant communication environments in most of the leagues, thus imposing a bias in the type of coordination architecture that can be deployed. Moreover, in the SPL data among teammates can be exchanged only once every second. There is, however, one league that hosts only centralized approaches. The SSL, in fact, has a single computer that receives sensory data streams; performs the computation; and deliberates collective and individual robot actions.

In the SSL we can observe the benefits of featuring a centralized coordination system where robots do not compete for shared resources; do not clash in ambiguous situations; and in general do not show the artifacts of a distributed approach where optimality in positioning and role-assignment is compromised to favor reactivity and individual behavior of a single player. Finally, we notice that in RoboCup, but also in other applications, centralized coordination architectures are deployed when robots can assume that the world is stationary and fully observable. Otherwise, it would be impossible to guarantee optimal behaviors even with a centralized strategy.

When considering extreme and dynamic environments, distributed architecture, are in general more robust to faulty communication; partial observability and non-stationary environments. The SPL, and in general all the humanoid leagues, are a clear example of such environments. Here, distributed coordination approaches are the most effective solution that researcher can resort to. Usually, a distributed coordination is achieved only exchanging local information among the teammates in order to reconstruct of global representation of the world state. Such information can have different format and might represent different concepts. The majority of distributed coordination systems exchanges either utility vectors or bid in auction-based methods in order to address dynamic-role assignment. Conversely, few approaches attempt to reconstruct a more sophisticated global world model by exchanging events that robots perceive locally and embedding them in a global (approximated) representation – which is updated iteratively. Finally, we notice that a common denominator of such techniques is that, distributed coordination systems are ready to recover from situation where a single unit might act individually and still trying to solve the task assigned to the team. With that firm in mind, existing approaches usually set a priority of the tasks to be complete. Then, the coordination system allocates roles in order to guarantee that the most important tasks are always active.

Cooperative vs. Adversarial. RoboCup is challenging from different points-of-views. In particular, given the structure of the problem and the environment, researchers are forced to investigate and find a solution to different problems at once. In this setting, for example, it is not possible to address separately multi-robot coordination and adversarial analysis.

Due to the individual leagues rules, an effective coordination is achieved differently across RoboCup competitions. Typically, coordination is achieved by balancing a strategic positioning of the different players within the field and dynamic role-assignment. In leagues where perception and hardware is not a limiting factor, coordination and cooperation also involves strategic setups, in-game schemes and multi-agent plays such as give-and-go. As observed in the previous sections, this is the case of simulated leagues and the SSL where perception is guaranteed to each agent; and specific to the latter scenario, computation is centralized.

Instead, we observe that dynamic role-assignment is constant in all leagues – where participants intuitively exploit the possibility to replace units without complications. Such a problem can be formalized as set of agents and a set of tasks, and the aim is to dynamically assign optimally each task to a particular agent. In the SPL, for example, role-assignment is achieved in a distributed fashion where each player takes individual actions but attempting to satisfy a team goal attempts to achieve a team. In this setting, approaches usually employ auction-based or utility-based solutions in order to trade-off robustness and reactiveness of the dynamic constraints of the environment. In the SSL, there are few approaches that implicitly assign all tasks/roles at once by means of a learning algorithm. They achieve good in-game performance but, such approaches are usually computationally inefficient and assume a reliable perception and localization. In fact, we notice that, the less these two assumptions hold, the simpler is the coordination strategy – which in extreme cases forces teams to implement static coordination approaches.

The adversarial nature of sports makes RoboCup an excellent testbed to promote research in the area of Adversarial MRS. However, there is only a subset of leagues that explicitly takes into account actions of the opponent teams, such as simulation leagues, SSL. But, there is an emergent trend in one of the humanoid leagues, i.e. SPL, that started to investigate how to react to opponent actions. In general, existing methodologies are coarsely categorized in behavior (and formation) classification and episodic reactive strategies. While the former involves long term strategies and re-positioning the robots in the field, as in the SSL, the latter shows more basic behavior that are reactions to movements of other players. Such strategies influence decisions of the robots and trigger local routines depending on the particular joint state of the players (both teammates and opponents). Also in this case, as for the coordination methodologies noisy-perceptions and hardware might become an obstacle that leads to techniques that promote robustness and compromise scalability.

Simulation vs. Hardware. We attempt to highlight which are the major factors that impact and influence the realization of effective multi-robot behaviors for physical agents in robotic soccer. The research conducted in simulated leagues supports the design and implementation of the most sophisticated coordination behaviors and collective strategies. In fact, given reliable perception and guaranteed computation power, simulated leagues have the privilege to only focus on behavior generation. From a different point-of-view, hardware oriented leagues like the Humanoid, historically focused their effort on behavior modeling for individual agents, limiting the environment to a pure adversarial setting. Recently, however, more robots have been added to the game, forcing teams to implement coordinated actions also for such platforms. Due to the constraints and size of such robots, current solutions to coordinated behaviors are basic and carefully handcrafted. But, the evolution of cooperative and collective behavior can only move forward, and as it happened for other leagues, researchers are experimenting with planning and machine learning approaches, moving away from ad-hoc modeling. We can conclude, that hardware is definitely one of the most critical factors in the development of coordination approaches. In fact to-date, heavy-platforms limit the research in collective behaviors and their deployment on the field.

Model-based vs. Learning-based. At the current stage, model-based approaches represent the most effective in-game choice for behavior generation and coordination due to their fast deployment and intelligibility. Learning-based methods, however, show the most promising results achieving more accurate and efficient solutions even if they suffer from high-computational demand, sample inefficient and they may result in brittle solutions. For these reasons, the use of pure learning-based approaches is limited and researchers tend to combine both paradigms to alleviate some of these issues. However, it is worth noticing that in leagues where high-computational support is guaranteed, and less complex platforms are used, learning approaches are more often deployed and have the chance to showcase their benefits. Examples of such environments are the Simulated Leagues and the Small-Size League. As a consequence, results shown by DL methods motivates teams to investigate learning-based approaches across all RoboCup leagues. In fact, during this survey, we encountered several learning methodologies which have been proposed to address and optimize particular in-game situations even in more complex and hardware-demanding leagues. But, to-date, such methodologies are deployed in an end-to-end fashion, and they are rather used to support planning-based solutions.

5 Conclusions

RoboCup competitions represent a exciting environment to foster novel research and excellent testbed to deploy and validate novel methodologies on physical agents. The research conducted within this context is fundamental develop autonomous end-to-end agents and already contributed to numerous applications outside the soccer field. We report that in recent years more sophisticated approaches have been proposed to optimize coordination and collective behavior at the decision-making level. However, it is worth noticing that not all the proposed approaches find a direct implementation in the actual competition. In fact, researcher tend to prefer planning-based approaches for in-game situations while learning-based methods seem to remain at a research phase. However, even if the solutions within the latter category are sample inefficient, not intelligible and may feature brittle behaviors, learning-based techniques show the most promising results improving performance in particular settings and generalizing to unknown scenarios.

Finally, we can summarize that the research conducted in robotics competitions is key for advancing and building effective teams of robotic agents. Competitions, in fact, represent a good balance between research and engineering of novel solutions that force the research community to develop techniques that can actually be implemented and can interact with the real world.

References

1. Adachi, Y., Ito, M., Naruse, T.: Classifying the strategies of an opponent team based on a sequence of actions in the robocup ssl. In: Behnke, S., Sheh, R., Sarıel,

S., Lee, D.D. (eds.) RoboCup 2016: Robot World Cup XX. pp. 109–120. Springer International Publishing, Cham (2017)
2. Adachi, Y., Ito, M., Naruse, T.: Online strategy clustering based on action sequences in robocupsoccer small size league. Robotics **8**(3), 58 (2019)
3. Akiyama, H., Tsuji, M., Aramaki, S.: Learning evaluation function for decision making of soccer agents using learning to rank. In: 2016 Joint 8th International Conference on Soft Computing and Intelligent Systems (SCIS) and 17th International Symposium on Advanced Intelligent Systems (ISIS). pp. 239–242 (2016). https://doi.org/10.1109/SCIS-ISIS.2016.0059
4. Antonioni, E., Suriani, V., Riccio, F., Nardi, D.: Game strategies for physical robot soccer players: A survey. IEEE Transactions on Games pp. 1–1 (2021). https://doi.org/10.1109/TG.2021.3075065
5. Bakkes, S.C., Spronck, P.H., Van Den Herik, H.J.: Opponent modelling for case-based adaptive game ai. Entertainment Computing **1**(1), 27–37 (2009)
6. Castelpietra, C., Iocchi, L., Nardi, D., Piaggio, M., Scalzo, A., Sgorbissa, A.: Communication and coordination among heterogeneous mid-size players: Art99. In: Robot Soccer World Cup. pp. 86–95. Springer (2000)
7. Catacora Ocana, J.M., Riccio, F., Capobianco, R., Nardi, D.: Cooperative multi-agent deep reinforcement learning in a 2 versus 2 free-kick task. In: Chalup, S., Niemueller, T., Suthakorn, J., Williams, M.A. (eds.) RoboCup 2019: Robot World Cup XXIII. pp. 44–57. Springer International Publishing, Cham (2019)
8. Dorri, A., Kanhere, S.S., Jurdak, R.: Multi-agent systems: A survey. IEEE Access **6**, 28573–28593 (2018). https://doi.org/10.1109/ACCESS.2018.2831228
9. Ghallab, M., Nau, D., Traverso, P.: Automated planning and acting. Cambridge University Press (2016)
10. Iglesias, J.A., Ledezma, A., Sanchis, A.: Opponent modeling in robocup soccer simulation. In: Fuentetaja Pizán, R., García Olaya, Á., Sesmero Lorente, M.P., Iglesias Martínez, J.A., Ledezma Espino, A. (eds.) Advances in Physical Agents. pp. 303–316. Springer International Publishing, Cham (2019)
11. Li, X., Chen, X.: Fuzzy inference based forecasting in soccer simulation 2d, the robocup 2015 soccer simulation 2d league champion team. In: Almeida, L., Ji, J., Steinbauer, G., Luke, S. (eds.) RoboCup 2015: Robot World Cup XIX. pp. 144–152. Springer International Publishing, Cham (2015)
12. Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 (2015)
13. Luger, G.F.: Artificial intelligence: structures and strategies for complex problem solving. Pearson education (2005)
14. MacAlpine, P., Barrera, F., Stone, P.: Positioning to win: A dynamic role assignment and formation positioning system. In: Workshops at the Twenty-Sixth AAAI Conference on Artificial Intelligence (2012)
15. Masterjohn, J.G., Polceanu, M., Jarrett, J., Seekircher, A., Buche, C., Visser, U.: Regression and mental models for decision making on robotic biped goalkeepers. In: Almeida, L., Ji, J., Steinbauer, G., Luke, S. (eds.) RoboCup 2015: Robot World Cup XIX. pp. 177–189. Springer International Publishing, Cham (2015)
16. Mendoza, J.P., Simmons, R., Veloso, M.: Online learning of robot soccer free kick plans using a bandit approach. In: Twenty-Sixth International Conference on Automated Planning and Scheduling (2016)
17. OpenAI: Openai five. `https://blog.openai.com/openai-five/` (2018)
18. Pierson, H.A., Gashler, M.S.: Deep learning in robotics: a review of recent research. Advanced Robotics **31**(16), 821–835 (2017)

19. Riccio, F., Borzi, E., Gemignani, G., Nardi, D.: Context-based coordination for a multi-robot soccer team. In: Almeida, L., Ji, J., Steinbauer, G., Luke, S. (eds.) RoboCup 2015: Robot World Cup XIX. pp. 276–289. Springer International Publishing, Cham (2015)
20. Riccio, F., Capobianco, R., Nardi, D.: Using monte carlo search with data aggregation to improve robot soccer policies. In: Behnke, S., Sheh, R., Sarıel, S., Lee, D.D. (eds.) RoboCup 2016: Robot World Cup XX. pp. 256–267. Springer International Publishing, Cham (2017)
21. Risler, M., von Stryk, O.: Formal behavior specification of multi-robot systems using hierarchical state machines in xabsl. In: AAMAS08-workshop on formal models and methods for multi-robot systems. pp. 12–16. Citeseer (2008)
22. Rizzi, C., Johnson, C.G., Vargas, P.A.: Fear learning for flexible decision making in robocup: A discussion. In: Akiyama, H., Obst, O., Sammut, C., Tonidandel, F. (eds.) RoboCup 2017: Robot World Cup XXI. pp. 59–70. Springer International Publishing, Cham (2018)
23. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of go with deep neural networks and tree search. Nature **529**, 484–503 (2016)
24. Spaan, M.T., Vlassis, N., Groen, F.C., et al.: High level coordination of agents based on multiagent markov decision processes with roles. In: IROS. vol. 2, pp. 66–73 (2002)
25. Suzuki, Y., Nakashima, T.: On the use of simulated future information for evaluating game situations. In: Chalup, S., Niemueller, T., Suthakorn, J., Williams, M.A. (eds.) RoboCup 2019: Robot World Cup XXIII. pp. 294–308. Springer International Publishing, Cham (2019)
26. Trevizan, F.W., Veloso, M.M.: Learning opponent's strategies in the robocup small size league. In: Proc. AAMAS. vol. 10. Citeseer (2010)
27. Watkinson, W.B., Camp, T.: Training a robocup striker agent via transferred reinforcement learning. In: Holz, D., Genter, K., Saad, M., von Stryk, O. (eds.) RoboCup 2018: Robot World Cup XXII. pp. 109–121. Springer International Publishing, Cham (2019)
28. Wurman, P.R., D'Andrea, R., Mountz, M.: Coordinating hundreds of cooperative, autonomous vehicles in warehouses. AI Magazine **29**(1), 9 (Mar 2008). https://doi.org/10.1609/aimag.v29i1.2082, `https://ojs.aaai.org/index.php/aimagazine/article/view/2082`
29. Yasui, K., Kobayashi, K., Murakami, K., Naruse, T.: Analyzing and learning an opponent's strategies in the robocup small size league. In: Behnke, S., Veloso, M., Visser, A., Xiong, R. (eds.) RoboCup 2013: Robot World Cup XVII. pp. 159–170. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
30. Zhou, Z.H., Yu, Y., Qian, C.: Evolutionary Learning: Advances in Theories and Algorithms. Springer (2019)
31. Ziparo, V.A., Iocchi, L., Nardi, D., Palamara, P.F., Costelha, H.: Petri net plans: a formal model for representation and execution of multi-robot plans. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1. pp. 79–86. International Foundation for Autonomous Agents and Multiagent Systems (2008)