The Standard Platform League

Eric Chown[1(✉)] and Michail G. Lagoudakis[2]

[1] Bowdoin College, Brunswick, ME, USA
echown@bowdoin.edu
[2] Technical University of Crete, Chania, Crete, Greece
lagoudakis@ece.tuc.gr

Abstract. The Standard Platform League is unique among RoboCup soccer leagues for its focus on software. Since all teams compete using the same hardware (a standard robotic platform), success is predicated on software quality, and the shared hardware makes quality judgments simpler and more objective. Growing out a league based on the Sony AIBO quadruped robots, the league has constantly evolved while moving ever closer to playing by human rules, and currently features the Aldebaran NAO humanoid robots. The hallmark of the league has been a focus on individual agents' skills, such as perception, localization, and motion, at the expense of more team-oriented skills, such as positioning and passing. The league has begun to address this deficiency with the creation of the Drop-in Challenge, where robots from multiple teams will work together. This new focus should force teams to work on multi-agent coordination in more abstract and general terms and promises to create fruitful new lines of research.

Keywords: RoboCup · SPL · NAO · Multi-agent cooperation · Machine learning

1 Introduction

The Standard Platform League (SPL) www.tzi.de/spl was referred to as the "Sony Quadruped Robot Football League" in the original rules published in 1998 [9] and was later renamed "Four-Legged League" (4LL). The original league adopted a standard hardware platform, the Sony AIBO, as a way to make the focus of the competition *software*, rather than hardware. The AIBO was a robust robot platform, which afforded teams the ability to do a significant amount of development and testing. AIBO was progressively replaced by the NAO robot between 2007 and 2009 and the league coined its current name. The latest edition of the league featured teams of five robots each, wearing red and blue jerseys, playing soccer autonomously on a 6×9 m field with yellow goals on both sides, using an orange street-hockey ball (Fig. 1).

The advantages of a standard platform are many. There are prominent teams in the league, which could not produce their own robots, but wish to use robots

R.A.C. Bianchi et al. (Eds.): RoboCup 2014, LNAI 8992, pp. 636–648, 2015.
DOI: 10.1007/978-3-319-18615-3_52

Fig. 1. SPL final game at RoboCup 2013 in Eidenhoven, The Netherlands. [10] (Color figure online)

and to work on software development. Further, the standardization of the hardware puts teams on theoretically equal footing, making the competition ultimately about the quality of the developed software. Since the teams all use the same robots, differences in the speed of motion, the accuracy of perception, the quality of localization, etc. can be directly attributed to software, rather than having to be disentangled from differences in hardware. This required focus on software has arguably pushed software development further in the SPL than in any other RoboCup humanoid league. This can be seen, for example, by the performance of Team DARWin which has won the KidSize class competition of the Humanoid League three years in a row using a code base developed for the SPL. That code base has since been widely cloned in the KidSize class. The focus on software can also be seen in the large number of research papers presented at the RoboCup Symposium that originate in the SPL.

2 History

2.1 The AIBO Years

The AIBO was an ideal platform for the original league, as the robots were relatively robust, and having four legs meant that their motion was relatively simple, as compared to two-legged robots, to control. Nevertheless, getting a small robot to play soccer presented a number of difficult problems in the early stages, especially with regard to vision and localization.

To help mitigate these problems the league used brightly-colored solid goals, in addition to a number of multi-colored beacons that ringed the field. While these unique landmarks helped to make localization much simpler, ironically they made vision something of a larger problem, because they increased the number of landmarks to be recognized, as well as the number of colors that a vision system needed to discern. In particular, AIBOs had trouble with the colors blue/skyblue, used in goals and beacons, and red/orange, used in balls and uniforms. Blue/skyblue colors were a problem, because they were close to so many colors that appear around a typical AIBO game (e.g. blue jeans worn by spectators), but also because the AIBO camera had a peculiar feature, called "chromatic distortion", that shifted parts of the image towards the blue end of the color spectrum. Red and orange colors were also a problem, because the balls used were close in color to the uniforms that the robots wore. And, since the robots were so low to the ground, it was easy to confuse their uniforms for balls.

The standard approach to vision was to use a color look-up table (LUT). Teams would take images from the competition venue and use them to "paint" a color table. The process was slow and tedious and prone to error. Because of this, attempts were made to improve the process with machine learning [1,8] Once a color table was built, most teams used a version of run-length encoding to extract blobs of the various colors that made up the landmarks around the field. Blobs were then examined and various sanity checks were run to see, if, for example, an orange blob was truly a ball and not part of a uniform. Notably, there was very little use of modern and general machine vision techniques, such as Hough and other transforms, mainly because of a desire to run at the highest frame rate possible. In addition, dedicated, bright, invariant, uniform lighting of the field was a necessity in those years for successful visual object recognition.

The fixed landmarks of the field (goals, beacons, lines) provided the cues for robot localization. Nearly every team ran some instantiation of the generic probabilistic Markov localization approach based on the idea of Bayesian filtering for state estimation. To update the state (position and orientation) of each robot, some kind of Particle Filter (Monte-Carlo Localization) or Kalman Filter (Gaussian Localization) or a combination of both was used. The advantage of Kalman Filters relied again on computational efficiency, as it was difficult to use enough particles to make Particle Filters work adequately on the limited AIBO processor. Similar techniques were also used for tracking the (moving) ball on the field. Over the years, the localization problem became harder, by altering the number and the quality of landmarks; the initially six multi-colored beacons around the field, became four, and finally just two, while the goals were simplified towards colored goalposts. In the last two editions of the Four-Legged League (2007 and 2008), only blue and yellow were used as landmark colors. Lines were not utilized much as landmarks due to their inherent perceptual ambiguity in contrast to unique landmarks and the additional complexity required in properly recognizing them, given that the robots themselves were also white-colored.

In terms of behaviors, the Four-Legged League was marked by a succession of improvements in individual skills and the development of basic behavioral frameworks, but very little obvious teamwork was evident. On an individual level, the AIBOs went from walking on their "paws", like normal dogs, to walking on what amounts to the forearms of their front legs. This made them much more stable and ultimately much faster and more agile. The dogs would kick the ball by putting their front legs up at a 90° angle and then forcefully bringing them down on the ball. This led to two important behavioral developments: the "grab" and the "dodge." Robots would grab the ball by throwing their chin on top of it and putting their front legs around it. This would stabilize the ball and put it under their control. It then naturally led to a behavior, where the robot could actually move from side to side, while holding the ball, to get to a better shooting angle, to the effect that attackers could actually dodge goalies and score goals by walking into the goal with the ball. The combination of these two strategies meant that the best teams could easily score multiple goals during a game. What really made this interesting, however, was that these behaviors could be acquired and optimized using machine learning.

The last few years of the Four-Legged League were marked by a significant rise in machine learning [2]. Teams discovered that their walks could easily be optimized by any of a variety of machine learning techniques [3,7]. Furthermore, individual behaviors, such as the grab [5], could be learned, as well as various aspects of color learning [8].

2.2 The NAO Years

After Sony ceased production on the AIBOs in 2006, the league began transitioning to a new robot platform. Following an open call for proposals presented at RoboCup 2007, the new standard platform was chosen to be the two-legged Aldebaran NAO robot. In 2008, SPL ran two competitions (AIBO and NAO) and in 2009 it transitioned completely to NAO. Moving to the biped NAO robot solved a number of problems that AIBOs had, but the change also created significant new problems.

The NAOs were able to benefit from ten years of improvements in robot hardware. For example, the NAOs has significantly better cameras than the AIBOs and faster processors. Among other things, this enabled the league to move to goal posts that were true posts. In addition beacons were no longer needed because the NAO's increased height enabled it to see more of the field than the AIBOs could. As a result, vision became simpler and ultimately localization improved, because the NAO's increased field of view made exploiting field lines possible. Eventually, these developments allowed to color both sets of goals posts uniformly yellow. While the new symmetric field created interesting new localization challenges, it also made vision even simpler by eliminating the blue color from the landmarks.

The reduction in colors (e.g. no more pinks on beacons, blues on goals or beacons) and landmarks (only one type of goal, no beacons, simpler uniforms), as well as better cameras and faster processors, has made more experimentation in vision possible. Teams, such as NAO Team HTWK and NAO Devils, have experimented with approaches to vision that do not use LUTs at all and instead infer the colors of the field of the objects by a combination of knowledge (e.g. most of the floor is green) and changes in luminance. Other teams have begun to move towards modern and general machine vision approaches, such as using Hough transforms to identify lines. Such approaches would not have been possible on the AIBO. What has not happened though is any sort of counterbalance to make vision a hard problem again. Ideas in this direction can easily be found in real soccer; the balls, for example, vary in color and patterns from game to game, as do team uniforms. Meanwhile, even though field lighting is not as carefully controlled as it was during the AIBO years, NAO games are still played indoors with fairly invariant ceiling, but no dedicated, lighting.

On the other hand, the league did work hard to counterbalance the improvements the NAOs brought to localization. The biggest change was moving to perfectly symmetric fields with regard to colors. No longer are there any unique landmarks that will tell a robot which direction it is facing. The current environment represents a significant contrast that to the early years, when there were two uniquely-colored goals, as well as six uniquely-colored beacons. Although the problem has gotten much harder, the underlying algorithms in localization have not changed significantly, besides incorporating field lines as landmarks and the need to rely on correct initialization. Probabilistic approaches based on Particle and Kalman Filters continue to dominate, but have been refined and extended with the addition of techniques, such as shared ball models and global (team) localization approaches, which in turn can help robots to disambiguate their location between the symmetric positions on opposite sides of the field.

By far, however, the biggest area of change and research, since the switch to the NAOs, is work on motion. Four-legged robots do not fall over and are relatively agile. Two-legged robots can fall and thus far are still far from being agile. In RoboCup 2009, the 2nd year of the NAO competition, the top three finishers were marked by the fact that each had developed their own walk engine (B-Human [6], Northern Bites [11], and NAO Devils [4]). Since then, there have been other teams who have developed very successful walks, e.g. rUNSWift and NAO Team HTWK, but the league has begun to converge on a single walk engine. Numerous teams now use some version of the B-Human engine. The reasons for this are pretty simple: a good motion system is still the dominant factor in robot soccer at the present. The best vision, localization, and behavior systems in the world cannot compete against a team that is simply faster to get to the ball and falls down less often.

While the motion systems used in the SPL now are vastly improved from what they were five years ago, the transition has cost the league in other areas. Robots that are slow, not agile, and prone to falling down do not make reliable teammates for coordinated activities, such as passing. Meanwhile, the robots are capable of kicking the ball more than the full-field length. This has led to a situation similar to the AIBO days, where there has been more emphasis on individual skills, such as directed kicking, than on coordination and passing. However, recent developments in the league point to a change towards this direction. One reason for this has been the development of "motion kicks", where the robots can kick the ball without first coming to a stop. Such kicks are not as powerful, but are much more useful against good teams, because they move the ball quickly and so it is less likely to be stolen by an opponent. Since motion kicks are less powerful, robots using them are not a threat to score from literally anywhere on the field and so moving the ball down the field in a series of kicks has become more important. As more teams use the B-Human, or similarly reliable, walk engine, walk speed will be cease to be the dominant differentiation factor between teams and team behaviors will become increasingly important. With precise kicking increasingly more prominent in the SPL, positioning of players and ball passing will start to become the primary factors that separate teams.

3 The Present

In an effort to spur development of better team play, the league has begun to experiment with "*drop-in*" games. Such games consist of teams of robots drawn randomly from the population of competing teams. To emphasize the importance of this competition, it is required that all qualified teams compete and some teams are invited to compete, even if they did not qualify for the main soccer competition. From the perspective of the league, there are two major challenges in running such a competition. The biggest challenge is *how to rank* the teams. Since one of the major goals of the competition is to increase interest in passing and teamwork, a blind, peer-review scoring system has been created with metrics that reward such teamwork activities. Nevertheless, judging a robot's quality in a given game is, at best, a difficult proposition. A poor decision may be the result of bad information provided by the teammates, but also a glimpse in the robot's own behavior system. The hope is that such variables will tend to even out over the course of a number of games. It is also clear that judging metrics will necessarily evolve over time to better meet the needs of the league.

A second challenge is *communication*. Drop-in games are not even possible, unless robots use a common, agreed upon communication protocol, so that they can share basic state information. To facilitate these games, the league has moved to a standard message packet format that each team is required to use in all games. The packet contains a number of standardized data fields for information, such as the robot's location and the location of the ball, and it also contains data fields that can be customized for any team's own needs, when playing normal games. The construction of the packet is not trivial, as the decisions have consequences for how well drop-in teams might perform. Early versions of the packet, for example, only contained basic information about the id of the robot sending the packet, where the robot was, and where the ball was. Such packets give no sense whatsoever of a robot's intentions. Is it going to the ball? Is it playing a particular position? And, so on. These things would need to be inferred by the robot's teammates. Ultimately, the league has chosen to include more information in the packet, including the robot's intended destination, where it is shooting to (or, is trying to shoot to), and a description of its intentions in terms of taking roles/positions, such as defender or keeper. Of course, the quality and even the truth of the information conveyed through these messages has to be considered by any robots receiving it.

There are many challenges involved in making such drop-in games work, and drop-in games could open rich areas of research for the league. For any given robot, for example, it is necessary to figure out which of its teammates it can trust. Some will be faster than others or have better localization systems. Some will communicate a ball location that isn't correct, etc. One change that this system may foster is the development of a more human-like position system. A theoretical advantage that robots have over humans is that they can all run the same code and have the same hardware, meaning that players can swap positions at any time. A single player might move from defense to striker and back to defense again, all in the space of a single point. If this is done seamlessly, it affords many advantages. In an ad hoc game, however, the amount of coordination to pull such a system off may not be achievable. Furthermore, in such scenarios robots might really have different capabilities. Therefore, it makes more sense for robots to adopt a human-like strategy, where players take a role in the beginning and stick with it. That way some of the communication that is lost by using a simple common protocol, can be implicitly gained back by knowing the basic roles of teammates. A human midfielder need not communicate explicitly with all of her teammates to know generally where they are and what they are doing, if she understands the basic principles of team soccer.

This year the league is also experimenting with a *coaching robot*. The idea comes from human soccer, where coaches provide general strategy for teams from the sideline. So, this year teams will be allowed to place a coaching robot on the sideline, where it can observe the game and communicate with the field players. The danger of such an addition, of course, is that the coach, who will have an fairly good view of the field, could provide solve problems for the individual players, such as which side of the field that they are on. To combat this, the league is experimenting with limiting the number and the style of communiques that the coach can provide, as well as putting a time delay on coach-player communications. The idea is that the coach should merely be communicating overarching strategies, not real-time state information. These communication limits are mirrored somewhat by a large drop in the amount of packets that field players are allowed to send and the requirement to use the new standard message format for communicating.

Meanwhile, solving the localization problems involved in a symmetric and fairly populated field remains an area of ongoing research. To this point no best practice has yet emerged, rather most teams seem to use a variety of heuristics. The most common has probably been to rely on teammates. For example a goalie can inform teammates when the ball is on its side. Other solutions have ranged from having a team's goalie make distinctive sounds to reveal the location of the home goal to using a variant of the SURF algorithm to try to build an on-the-fly map of the field's surroundings.

In terms of the mechanics of soccer, the league has once again progressed back to the point where the games were at the end of the AIBO era; there are five players on each side and the field has grown to be 9×6 m. Typically, the league sticks with a given field size and number of players for approximately two to three years, so it is likely that the league will look to use larger fields with more players in the next year. While this growth represents important progress towards the ideal of playing on human fields, it brings its own set of challenges. Some of these are purely practical—not many labs are large enough to support a full-sized field and, of course, more robots require commensurate amounts of additional money. There is a further complication as well having to do with how much teams can run practice games.

In principle, one of the advantages SPL has over all other RoboCup soccer leagues is the common standard robot platform. Among other things, this means that it should be possible to run regular practice games between players of the same team and even to play practice games against other teams, if their code (or their executables) are available. This actually happened a fair amount of times in the last few AIBO years. For example, the Northern Bites were able to compete remotely in the RoboCup German Open competition in 2008. In turn, the German Team provided the Northern Bites with executable versions of their code so that the Northern Bites could play practice games against the German Team in their lab. The advantages of such arrangements are myriad, but boil down to the fact that the way to get better at anything is to practice. In RoboCup, practice games provide essential data and opportunities for debugging, considering the variation offered by playing against other teams. In some ways there are more opportunities for these kinds of practice games now than ever, as several teams release their code on a yearly basis. However, as the league rules, including field dimensions and number of players, change, the amount of work necessary to get one of these releases running, obeying the updated rules, quickly becomes prohibitive. Further, even with the current settings, running a full practice game requires a total of 10 healthy robots, a number that very few teams can afford.

4 The Future

The move to two-legged robots was a case of the league needing to take a number of steps backwards in terms of quality of play in order to make some very large steps forward possible. The difficulty of walking stably and quickly was clearly the dominant focus of the SPL for several years after the switch. At this point, the low-hanging fruit of what the NAOs can do has more or less been reached and many of the remaining limitations can now be attributed to the hardware. On the one hand, such limitations are a roadblock to further development in terms of motion, but, on the other, it means that the standings of the league will not be so dependent on who has the best walk engine, which ought to force developments in other areas. The biggest area with the most potential for future developement is almost certainly team play, such as team formations, role assignment, player positioning, and ball passing. In order for a team to develop a successful passing strategy, there must be a combination of ability and need. On the ability side, robots must know where their teammates, as well as their opponents, are and be able to kick the ball with a fair amount of precision. Receiving robots must be agile enough to stop and control the ball. On the need side, it must be the case that it is better for a robot to pass the ball than to simply make a long kick towards the opponent goal or try to dribble the ball upfield on its own.

In terms of ability, many of the pieces are in place. Robots in the league localize well enough that most robots very reliably know where their teammates are with a relatively small amount of error. While opponent localization is a much harder problem, many teams can now reliably identify opposing robots in their field of view and are actively working on incorporating them into a shared model. Kicking the ball accurately remains difficult, because of the shapes of the robots' feet, limitations on the movements of their joints, and differences in field conditions. Nevertheless, robots are able to kick the ball with some accuracy over short distances. The biggest limitation seems to lay in the receiving robots. Even the best humans do not always make perfect passes, but human receivers do quickly adjust to imperfect passes. Humans can quickly shift their feet, their motion path, or whatever is necessary in order to receive and control a pass. Robots in the SPL are simply not capable of doing much of that yet. Simple passing, not nearly as fluid as it is for humans, still remains to be seen.

In terms of need, a major reason there is so little passing seems to be that there is not enough payoff to try it very often. A pass should provide the passing team with a clear advantage over not passing. In RoboCup, a player near the ball has three basic options: shoot, dribble, or pass; right now, passing is the third choice for most teams. This mainly has to do with the fact that shooting is such a high-reward, low-risk option. If a player is aligned toward the goal, there is little reason not to shoot. Players can easily kick the ball hard enough to score from anywhere on the field. Even if the goalie stops such a shot, the ball has gone to a place where the shooting team is at an offensive advantage, particularly since there is no offside rule and many teams position players near the opposing goal. The only disadvantage of such long shots is the possibility that the ball goes out of bounds; a rule introduced in recent years to discourage random shooting dictates that, in such a case, the ball is replaced in the field one meter behind the kicking robot. On top of that, robots that take the time to stop and set their feet to kick over long distances risk having the ball stolen by opponents. Given the lack of precise kicking, the next best option is to dribble, possibly using motion kicks, which are less powerful, but much quicker and almost akin to dribbling. Some addition to the rules may be required to discourage continuous dribbling and encourage work on the next option (passing). However, this may naturally occur, as the number of players and the field dimensions change.

One long-term goal of the SPL has been to be the first humanoid robot league to play 11 versus 11 players' games. With more players on the field necessarily comes an increased premium on teamwork. In the interim the league will continue to push towards games that are more realistic. Changes expected to occur in the near term include changing the goal colors to white and letting teams wear custom jerseys. Both will present new challenges for league vision systems, and both are necessary to the evolution of the league, but neither represents the kind of challenge that will significantly advance the quality of play in the league.

The very fact that RoboCup is a competition encourages teams to stick to "best practices" at the expense of breaking ground on radically new approaches. In recent years this has been reinforced, as the best teams have once again begun to regularly release their code and numerous other teams have used that code as the basis for their systems. On the one hand, this is desirable, as it pushes the overall quality of play higher, allows new teams to get a working system up and running, and forces the best teams to keep improving. On the other hand, many teams make little or no effort to truly push the boundaries of robot soccer.

Going forward the league requires balancing several things. The field needs to get bigger and more robots need to play at a time. On the other hand, if the trend is pushed too far/fast it will eliminate the ability of some teams to practice or even compete. In the past the league has discussed facing the issue by forcing the merger of teams into larger regional teams. The current move towards playing drop-in games, however, seems to be a more pragmatic solution and one that has payoffs in other dimensions. With drop-in games established, each research team no longer necessarily needs to have more than a few robots. The downside is that practice games are practically eliminated. In the long term, it may raise the need for more local competitions. Currently these local competitions are mainly seen as a tune-up for the main RoboCup event, but perhaps with more incentives in place they could be more useful and could even include remote participation. Ultimately, robots are so useful as a tool of science, precisely because they act in the real world. The more data that we can get from such actions the better.

The Standard Platform League needs to work to ensure that it leverages the qualities that make it unique among RoboCup leagues as much as possible. The Drop-in challenge is a strong, positive step in that direction. With the increased emphasis on this portion of the competition teams should begin to think beyond how to engineer the best performance for their team and more about how to create flexible agents capable of thriving under a wider variety of circumstances and with changing teammates. In turn, this could lead the league to explore research areas that it would not have considered otherwise.

To provide a specific example of how such research exploration can be realized, we note that any kind of communication between players in the SPL currently relies solely on the availability of a reliable wireless network. This is, however, an artificial communication channel, missing in real soccer games, whereby human players exchange and share information using audiovisual signals they generate. In addition, experience from the most recent RoboCup competitions indicates that even the current state-of-the-art technology in wireless communication is insufficient to support reliably the needs of SPL games. In the recent 2012 and 2013 SPL games many robots could not even connect to the network, players could not receive game state from the game controller, data packets were constantly missing, teammates could not communicate to coordinate, and eventually most efforts for successful coordination among team players were gone. Since communication is obviously necessary to advance teamwork, the key question that the league may need to address in the future is the following: How could robots exchange and share information, if there was no wireless or any other

Fig. 2. SPL participants at RoboCup 2013 (picture credit: Brad Hall).

kind of network in the SPL field? Could they use the physical world? Could they generate visual and/or auditory cues that could be perceived by other robots and thus convey information in a more natural and human-like way?

As a concrete example towards audiovisual communication, consider the problem of generating an intuitive visual signal for sharing information about the most important element in a soccer game: the current position of the ball in the field A robot having direct visual observation of the ball can assist its teammates by generating a visual signal that reveals the current position of the ball. The most natural and intuitive visual signal one could imagine is to fully extend one of the two arms, so that it points directly to the ball. Another robot could visually perceive this signal through its camera and use the vector defined by the first robot's arm to pinpoint the position of the ball. Note that this kind of information sharing about the position of the ball does not necessarily require any kind of localization of the two robots in the field or relative to each other. The perceiving robot could simply "follow" visually the line indicated by the pointing robot to locate the ball. Furthermore, such a visual signal could be used for directing attention to any desired position or object in the field; a lead player could direct a teammate towards a certain position in the field and a future robot referee could easily point out the player that was penalized or the place where the ball went out of bounds! Considering also the new addition of a robot coach on the sideline, images of real human coaches would come to mind, if one watched the robot coach pointing to positions and players, possibly "shouting" instructions at the same time.

The elimination of the (artificial) wireless communication network will open up opportunities for researching several human-like behaviors, which may not be strictly restricted to soccer. Teams will have to develop their own (or adopt existing) sign languages to convey information between their players, possibly taking into account the fact that opponents are also "watching" and information may have to be encoded. Robot perception will not be limited anymore on just correct object recognition (ball, goals, lines, players, etc.), but will also have to succeed in correctly recognizing postures, signals, gestures, sounds, even complete “oral” sentences in a human or even in an (artificial) robot language. Needless to say that all competing players will have to obey referee instructions, signaled with the sounds of a whistle, possibly complimented with details provided “orally”. Under the described setting, not only will human spectators be naturally engaged into the game play strategies of each team, but also the technology developed could be further exploited in numerous human-robot and robot-robot interaction situations beyond robot soccer.

5 Conclusion

The SPL remains one of the most vibrant leagues in RoboCup (Fig. 2). It has always been a leader in advancing humanoid robot soccer in increasingly realistic directions in terms of the field dimensions, the number of robots on the field, and the setup of the field. In turn, the emphasis on software has led the league to develop and refine techniques in motion, vision, localization, communication, behavior, and coordination, but also to experiment with applying machine learning and optimization techniques to a range of tasks. In the present the league is continuing to lead as it experiments with the drop-in challenge and with coaching robots. In the future, it will continue to push the boundaries of what robots can do on a soccer field.

References

1. Bruce, J., Balch, T., Veloso, M.: Fast and inexpensive color image segmentation for interactive robots. In: Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol. 3, pp. 2061–2066 (2000)
2. Chalup, S., Much, C., Quinlan, M.: Machine learning with AIBO robots in the four-legged league of RoboCup. IEEE Trans. Syst. Man Cybern. Part C Appl. Rev. **37**(3), 297–310 (2007)
3. Chernova, S., Veloso, M.: An evolutionary approach to gait learning for four-legged robots. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), vol. 3, pp. 2562–2567 (2004)
4. Czarnetzki, S., Kerner, S., Urbann, O.: Observer-based dynamic walking control for biped robots. Robots Auton. Syst. **57**(8), 839–845 (2009)
5. Fidelman, P., Stone, P.: The chin pinch: a case study in skill learning on a legged robot. In: Lakemeyer, G., Sklar, E., Sorrenti, D.G., Takahashi, T. (eds.) RoboCup 2006: Robot Soccer World Cup X. LNCS (LNAI), vol. 4434, pp. 59–71. Springer, Heidelberg (2007)
6. Graf, C., Härtl, A., Röfer, T., Laue, T.: A robust closed-loop gait for the standard platform league humanoid. In: Proceedings of the 4th Workshop on Humanoid Soccer Robots at the IEEE-RAS International Conference on Humanoid Robots, pp. 30–37 (2009)
7. Kohl, N., Stone, P.: Policy gradient reinforcement learning for fast quadrupedal locomotion. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA). vol. 3, pp. 2619–2624 (2004)

8. Quinlan, M., Chalup, S., Middleton, R.: Application of SVMs in colour classification and collision detection with AIBO robots. In: Advances in Neural Information Processing Systems (NIPS) (2004)
9. RoboCup: Sony Quadruped Robot Football League Rule Book (1998). www.tzi.de/spl/pub/Website/History/Rules1998.pdf
10. Snapshot from BotSport.TV. www.youtube.com/user/BotSportTV (2013)
11. Strom, J., Slavov, G., Chown, E.: Omnidirectional walking using ZMP and preview control for the NAO humanoid robot. In: Baltes, J., Lagoudakis, M.G., Naruse, T., Ghidary, S.S. (eds.) RoboCup 2009. LNCS, vol. 5949, pp. 378–389. Springer, Heidelberg (2010)