Participating in Autonomous Robot Competitions: Experiences from a Robot Soccer Team

Martin Lauer
Institute of Measurement and Control
University of Karlsruhe
lauer@mrt.uka.de

Martin Riedmiller
Institute of Computer Science
University of Freiburg
riedmiller@informatik.uni-freiburg.de

Abstract

This paper describes the development of a RoboCup team during six years of participation in RoboCup. From the point of view of a participating team it discusses the experiences made, the scientific impact, and the correlation between effort, tournament success, and scientific achievements.

1 Introduction

During the last 15 years scientific competitions have emerged as a new platform for scientific comparison and exchange. In contrast to classical forms of scientific events like conferences and workshops competitions do not only allow to present new algorithms and theoretical results but also allow direct comparison between different approaches for complex problems.

One of the first competitive events is RoboCup, an international competition of autonomous soccer playing robots [Kitano *et al.*, 1997] which has started in 1996 and offers yearly competitions (see also fig. 1). Meanwhile, RoboCup has become a very large event with competitions in different leagues of robot soccer, rescue robots, human-robot interaction, and educational robots. The leagues differ in the robot size, capabilities, and the tasks to be done. In this paper we want to focus on the so-called MiddleSizeLeague, in which completely autonomous, wheeled robots of 80cm height play soccer. The games are played between two teams of five robots each and take 30 minutes. The rules of play are adapted from human soccer. Each team creates its own robots so that the robots differ in mechanical design and control software.

With the team *Brainstormers Tribots* we participated in the MiddleSizeLeague for six years between 2003 and 2008 which gives us the possibility to pass in review the development of the RoboCup MiddleSizeLeague competition and to discuss its scientific impact and its benefits. Due to the topic of this paper some arguments presented throughout the text are personal assessments. The paper aims at describing our experiments and helping researchers interested in participating in competitive scientific events to assess the risks and possible rewards. Since we want to focus on the relationship between competitive scientific events and research we will not discuss the impact of robot soccer on education, publicity, and industrial cooperations which constitutes another important aspect of RoboCup.

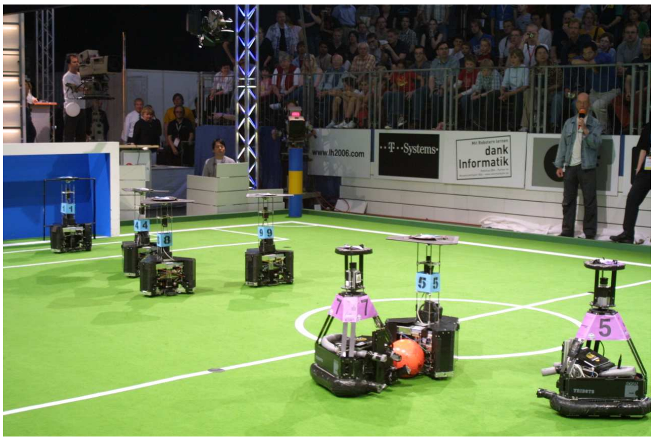

Figure 1: Picture from the 2006 finals in the RoboCup MiddleSizeLeague between the teams *Brainstormers Tribots* (magenta) and *CoPS* (blue).

The paper is organized as follows. In section 2 we describe the general development of our RoboCup team, of the robot hardware and the control software. Subsequently, we analyze in section 3 which scientific impact this project had for our research group and in which way the scientific work flow differs between classical research and research in projects like robot soccer. Finally, in section 4 we discuss the benefit of robot soccer activities from a subjective perspective and make proposals for an improvement of the ratio between scientific impact and effort.

2 The Team Brainstormers Tribots

2.1 Team Development

The RoboCup team *Brainstormers Tribots* was initiated in the year 2002 in the research group of professor Martin Riedmiller at Dortmund University. At this time the research group included only professor Riedmiller, three Ph.D. students, and some master students. Already four years earlier Martin Riedmiller has founded a RoboCup simulation league team named *Brainstormers* which was participation in RoboCup with big success and which was further developed

year	2003	2004	2005	2006	2007	2008
German Open (2006: Dutch Open)	quarter final	first	first	first	first	third
RoboCup championship	second round	quarter final	x	first	first	third
Technical Challenge	-	-	x	first	first	first

Table 1: Success in the official RoboCup tournaments. "-" refers to an unknown rank not among the first three, "x" refers to not participated.

in parallel to the MiddleSizeLeague team.

Based on a hardware concept for mobile robots developed in a master's thesis a team of 13 master students started to create the MiddleSizeLeague team within half a year. The main work including the design of the hardware and software was done by the students while the researchers acted as advisors and supervisors. At the German Open 2003 – a European RoboCup championship – the team was participating the first time and successfully reached the quarter finals where it dropped out. The participation at the RoboCup world championship followed.

The successes of the first year motivated a core group of volunteers to continue the work on the robot hardware and software. The core group included three Ph.D. students and five master students. From this point on the project turned from a practical course into a research project and scientific questions became more and more important.

In 2004, the team could win the German Open for the first time while it dropped out of the world championships during the quarter finals. The subsequent year 2005 was characterized by continuous work to improve the performance and reliability of the robots. A participation at the world championships was denied by the team members to save manpower and get the opportunity to do scientific work also outside of robot soccer.

The RoboCup world championships 2006 were held in Germany so that the team was highly motivated and the year became very successful winning the Dutch Open, the RoboCup world championships and the so-called *technical challenge*, a competition for the best scientific and technical development of the year in the MiddleSizeLeague. In 2007, the success of 2006 was repeated while in 2008 the team became third at the German Open and third at the world championship while the team was winning the Technical Challenge a third time. Table 1 shows the major results of our team in the RoboCup competitions.

During the six years between 2003 and 2008 there were only little changes in the core group of researchers and students. Even after finishing their master's thesis and working in industry some students continued to participate in the team. Additional students joined the team for one or two years.

2.2 Hardware Development

Being mainly interested in machine learning techniques and neuroinformatics the research group of Martin Riedmiller was located at the computer science department of Dortmund University until 2003 and at the Institute of Cognitive Science at the University of Osnabrück later on. Hence, all team member were computer scientists or students in cognitive science. Expert knowledge in mechanical and electrical engineering was scarce. Luckily, our team started its development when all teams had to rebuild their systems integrating holonomic drives [Pin and Killough, 1994] and catadioptric camera sensors [Baker and Nayar, 1998]. Thus, all teams had to gain experiences with the new techniques. Using standard components as much as possible like laptops, standard industrial cameras, and motor controllers, we were able to create a robot on the state of the art [Hafner, 2003].

The simplicity of our robot design, although ugly to watch, turned out to be a big advantage since unforeseen problems could be solved easily using adhesives or tape. By trial and error we found easy hardware solutions for problems like controlling the ball using foam fingers or implementing chip kicks using a warped kicking device. Some of these solution have been copied by many other teams later on.

The first four robots built in 2002/2003 exhibited some minor shortcomings that caused us to rebuild five new ones in 2004 which replaced the first generation (see fig. 2). However, only minor changes of the design of the robots were made to improve usability and robot velocity. While we had a lot of damages and breakdowns in the first two years the hardware became more and more robust after fixing the problems. Many problems were caused unexpectedly by wires and plugs which lost contact due to the vibrations of the robot or by electric charging of the robots.

Until 2007 the robots built in 2004 could compete against the newly built robots of other teams. However, the further development in the design of soccer robots led to robots with stronger kicking devices, faster motors, cameras with higher resolution and computers with much more computing power. Furthermore, it became difficult to get spare parts and to keep the robots running so that we were forced to develop new robots. Some parts for a new robot generation have been used in 2008 but they exhibited a lot of teething problems.

2.3 Software Development

In contrast to the mechanical design of the robots, the software development was much more in the focus of our team. However, being a research group in the domain of machine learning we did not have much experience with software development for autonomous, mobile, real-time systems. Therefore, the first control software developed in 2003 was a kludge and building a long-term research project on it was not possible. After an interim software solution used in 2004 we developed a well-structured, modular software framework which is used since 2005. To guarantee that the software remains maintainable it was necessary to clean up the software occasionally and to replace kludge, which was usually generated during the tournaments, by well-structured code. Here, the existence of a core group was beneficial.

Figure 2: Robot of the team *Brainstormers Tribots* in the year 2007. From top to bottom the picture shows the omnidirectional camera, an additional perspective camera, the control laptop, the kicking device, and the chassis.

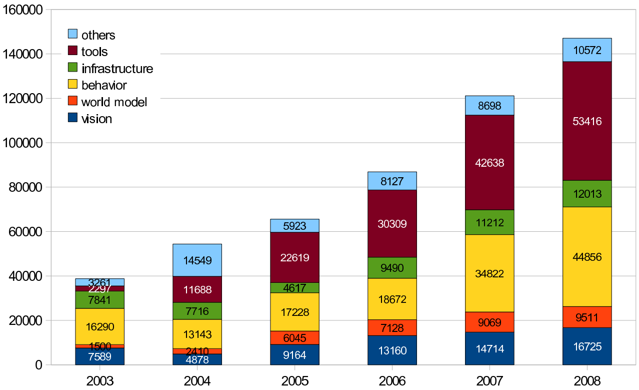

Figure 3: Lines of code of the Tribots control software between 2003 and 2008. The different colors refer to different parts of the software.

The development of the software in terms of lines of code is depicted in figure 3. The total amount of code grew from 39,000 lines in 2003 to 147,000 lines in 2008, i.e. the amount of code grew on average by 30% per year. Interestingly, the chart shows large differences between different parts of the software. We distinguish between five areas, the vision subsystem that implements the image processing and object recognition algorithms, the world model that collects sensory information over time and calculates a consistent geometrical and dynamical model of the environment, the behavior component that creates the robot behavior, and the infrastructural code including the software framework, control loop, communication with the motor controller and with teammates. Additionally, some separate tools are implemented for software development, debugging, visualization, and calibration of the robot.

The different growth rates of the software components reflect the varying activities of the team. While during the first years the development of the basic infrastructure, a reliable vision system and a consistent world model constituted the key activities, the further development of the robot behavior became much more important during the second three years. Surprisingly, the development of tools for visualization and calibration is the area with the largest growth rate. In the year 2008 more than a third of all lines of code contributed to these tools. On the one hand this reflects the large effort that is necessary to implement user-friendly tools and graphical user interfaces, on the other hand it shows that complex tools are necessary to monitor and analyze complex robot behavior.

3 Scientific Impact and Innovations

3.1 Scientific Workflow

The original intention of our research group starting to build an autonomous robot was to have available a mobile robot for experiments in reinforcement learning [Gabel *et al.*, 2006]. However, the necessity to provide the robot with a visual perception, a representation of its environment, and an easy to use software framework quickly opened up large field of research like vision, sensor fusion, cognitive and agent architectures, and multi agent coordination. Moreover, the development and implementation of reliable methods for these domains took several years so that we had to wait three years until we could start with reinforcement learning experiments on the robot.

However, the need to develop algorithms in the aforementioned areas created a new research focus for our group. We want to exemplify this phenomenon in the following for the domain of perception and sensor fusion. A soccer robot needs to know its environment to be able to interact and behave reasonably. Important variables are the own pose, the own velocity, the position and movement of the ball, and the position and movement of its teammates and opponents. For higher levels of reasoning it must also be able to perceive relationships between objects like *robot A dribbles the ball* or *robot B plays a pass to robot C*. However, the latter relationships might also be deduced from the former observations. Obviously, the better the variables mentioned can be estimated the better the behavior can be.

The development of these capabilities of perception, sensor fusion, and representation proceeded step-by-step. In the first year, a simple approach of self-localization was developed based on a particle filter [Thrun, 2002] and perception of colored areas on the back wall of the goals. The perception of the ball and obstacles was also based on color segmentation. For simplicity, we did not estimate the velocity of the ball and the obstacles. Certainly, this first attempt did not work reliably and the precision obtained was low. This approach was below state-of-the-art.

Consequently, we developed a second approach for self-localization combining a particle filter with recognition of the white field markings. This approach was on the cusp of state-of-the-art so that we could publish a first workshop paper [Merke *et al.*, 2004]. In the subsequent years we further improved self-localization creating a new algorithm [Lauer *et al.*, 2005a]. Beside self localization we also developed methods for the estimation of the ball position and movement [Lauer *et al.*, 2005b; 2006], for the ego motion of the

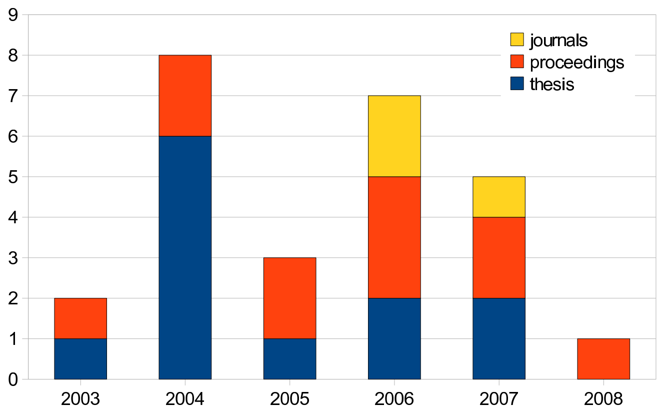

Figure 4: Publications between 2003 and 2008 not including team description papers. *Thesis* refer to Bachelor's and Master's thesis.

robot, and for the recognition of situations in which the robot is blocked by other robots or obstacles [Lauer, 2006]. As result of this work we did not only obtain a reliable implementation of a geometric and dynamic model of the environment but by putting together our research papers we also obtained a theory of robot perception, at least for the domain of a soccer field.

More than this, RoboCup is a dynamic benchmark that further develops year by year. Rule changes and improvements of the opponent teams increase the complexity continuously and require ongoing further developments. For instance in 2005 some teams started to execute chip-kicks while before the ball was almost always remaining on the ground. Hence, we had to further develop the visual perception of our robots replacing the monocular vision by a binocular approach combining the catadioptric camera with an additional perspective camera. Again, the needs of soccer robots opened a new field of research [Voigtländer *et al.*, 2007; Lauer *et al.*, 2009] which we potentially would have never entered without robot soccer.

The example described above shows how developments in competitive events like RoboCup are typically made. Most of them are driven by the necessities of the competition. Year by year new building blocks are added and tested in the tournaments. Usually, the presentation on a conference or in a journal follows with a delay of one year. This way of working is very different from the classical scientific workflow in which the experimental evaluation under real-world conditions is done much later, if at all an evaluation under real-world conditions is done. On the other hand, publications from robot soccer often suffer from a lack of theoretical depth since the development is not driven by theoretical considerations but by practical requirements and the step to lift heuristics and ad-hoc approaches to a theoretical level is not done so that papers submitted for publication are often rejected.

From our own experience gathered during the last six years we know that there are a lot of developments in robot soccer that have not been published yet due to concentration of manpower onto the tournament results. Unfortunately, this also applies for our own team. Figure 4 shows the number of publications that have been created by our team.

1.	learning to dribble with reinforcement learning
2.	visual recognition of red cards, player exchanges, and referee gestures
3.	kicking and dribbling device
4.	omnidirectional camera calibration and situation dependent role assignment
5.	fault detection and fault tolerant behavior
6.	playing with a white ball
7.	software framework for behavior based agents
8.	cooperative stereo vision
9.	FPGA based vision
10.	kicking device
11.	team description

Table 2: List of contributions of the free technical challenge in the MiddleSizeLeague 2007

3.2 Technical Challenges

Beside the robot soccer tournaments so-called *Technical Challenges* have been established in some leagues as another form of comparison and scientific exchange. In these competitions special tasks are defined that have to be executed by the teams. Often, these tasks focus on problems that are not yet necessary in the present tournament but which will become important in future. E. g. in the RoboCup MiddleSizeLeague during the last three years these challenges were facing pass playing, self localization without color information, and visual perception of arbitrarily colored balls.

A second form of technical challenge are so-called *free challenges* in which each team can present results of its scientific work and technical further developments. Table 2 exemplifies the work presented by the MiddleSizeLeague teams in the 2007 tournament. It becomes clear that the teams have different research foci like machine learning, vision, mechanical engineering, fault tolerant systems, multi-agent-systems, etc. Although parts of the work presented there were impressing the technical challenges are not visible from outside RoboCup, even from outside of the MiddleSizeLeague. A presentation to the community outside of the own league is missing completely, even a simple listing of topics is not available so that the teams do not spend much effort in preparing their contributions.

3.3 Innovations of Team Brainstormers Tribots

At this point let us take a look on the innovations that our own team contributed during the last six years. We want to distinguish between scientific contributions which have the potential to be published in scientific literature or already have been published, and technical improvements which are related to the soccer playing task and which are of no interest outside robot soccer.

The first group comprises the development of a consistent description of the environment described in section 3.1 and the development of a stereo camera system combining a catadioptric video sensor with a perspective sensor. Furthermore, we developed techniques in reinforcement learning that enabled the robot to learn autonomously tasks like intercepting a ball [Müller *et al.*, 2007] and dribbling a ball [Riedmiller *et al.*, 2008] which is closely related to our original motivation to build soccer playing robots. The techniques developed for soccer robots could be transferred to other areas like learning the steering controller of an autonomous car [Riedmiller *et al.*, 2007] or solving multi-agent scheduling problems [Gabel and Riedmiller, 2007]. Additionally, we created a modular cognitive architecture for autonomous robots and showed that this architecture is easily transferable to soccer robots of other teams.

The innovations which are closely related to soccer robots are the development of dribble control approaches, chip-kicking, cooperative play in the soccer domain, a very successful attacking and defending strategy, pass playing, and a development framework for soccer robots.

3.4 Developments in the RoboCup MiddleSizeLeague

A competition like RoboCup does never reach a final research goal but the research challenges are defined by the capabilities of opponent teams and further develop year by year. E.g. in 2002 one team created a kicking device that was much harder than all devices known so far so that it could win against almost all opponents. This innovation induced better defense strategies and faster perception algorithms to be able to defend one's own goal. However, to play successfully against teams with strong defenses capabilities like dribbling and chip kicking were invented. Then again chip-kicking induced the development of stereo camera systems.

The future developments cannot be forseen exactly. The development of pass playing is claimed since many years. However, although some teams like ours are able to play passes it does not play a major role during the tournament up to now since the advantage of pass playing is not that large as the advantage of precise chip-kicking or fast dribbling. For the next years the authors expect further developments in the precision of sensing and acting, the visual discrimination of teammates and opponents, better cooperative play including pass playing, adaptive player behavior and automated game analysis.

Beside the competitive argument rule changes force the teams to adapt their software and hardware to new requirements. The general idea behind rule changes is to make RoboCup rules more similar to human soccer rules. E.g. the walls around the fields were removed in 2002/2003, the field size grew from $5 \times 8m$ (2002) to $12 \times 18m$ (2008). Furthermore, set-plays were introduced in 2005, the playing time was extended from 20 minutes to 30 minutes (2007), and the number of robots increased. Moreover, since 2008 the goals are not color coded anymore and for future, the special ball color will also be removed. A list of the major changes between 2002 and 2008 is given in table 3.

All of these changes caused adaptations of the teams. E.g. the removal of walls required a complete change of the self-localization approaches since before most teams used laser scanners to recognize the walls while since 2003 vision based approaches are needed. The increase of field size also caused changes on the video sensors and perception approaches to

	2002	2008
field size	$5 \times 8m$	$12 \times 18m$
game duration	20 min	30 min
field boundary	walls/poles	white lines & security boundary
game interrupts	none	set-plays similar to human soccer
color coding	goals, corner poles, ball, field markings, team labels, robots	ball, field markings, team labels, robots
lighting	artificial, homogeneous	mixed artificial/ sunlight, varying over time and place
sensors	laser range scanner, perspective cameras	catadioptric cameras, additional perspective cameras
drive	differential drive	holonomic drive
robot velocity	$1 - 1.5\frac{m}{s}$	$2 - 3\frac{m}{s}$
kick strength	$2 - 3\frac{m}{s}$	$3 - 7\frac{m}{s}$
chip kicks	none	often, up to $2m$ height
dribbling	none	often
pass playing	none	sometimes
strategy	ego-centered	cooperative

Table 3: Major changes in the rules of the RoboCup MiddleSizeLeague between 2002 and 2008 (upper half) and changes in the dynamics and strategy of the robots (lower half)

be able to cover a larger part of the environment. On the other hand, the extension of the fields allowed to play with higher dynamics, to dribble the ball and to play passes which is almost impossible on small fields.

4 How Robot Soccer Contributes to Science: a Subjective Summary

4.1 Present Situation

After six years of active participation let us allow to give a subjective summary of activities in robot soccer and their impact onto the scientific development. As we have described in section 3.1 participating at competitive events in research induces a very different scientific workflow driven by the real-world application rather than by theoretical findings. Therefore, this form has been criticized in past.

However, in areas like artificial intelligence and autonomous robots the empirical evaluation of approaches is very important and cannot be replaced by a theoretical analysis alone since the environment with which autonomous robots have to interact is usually hard to specify mathematically. This basic idea is refelected by the *embodiment*-hypothesis [Pfeifer and Bongard, 2007] which claims that intelligence requires physical existence. This hypothesis has been discussed deeply during recent years and many arguments from cognitive science support it. Hence, the study of artificial intelligence and autonomous cognitive systems must investigate physical systems and their interaction with the environment. Robot soccer is one among many possible benchmarks exhibiting high dynamics, collaboration, and physical interaction with the environment. Thus, it is helpful to study aspects of perception, multi-agent-collaboration, and behavior generation, among others.

But not only the workflow is different from classical research but also the way of thinking in systems rather than in small, modular problems. E. g. in the domain of machine learning classification algorithms assume a very abstract problem, i.e. some patterns in terms of mathematical vectors have to be partitioned into two classes. From which process these patterns come and which consequences a wrong result has is of no interest at this level of abstraction. While the abstract view is appropriate to compare the performance of isolated algorithms among each other it is of limited use to measure whether such an algorithm is appropriate to solve a real-world problem.

In robot soccer, the focus is completely different since that team wins which has the best overall solution. The combination of very different techniques to a complex system is more important than the optimization of each individual module. Since the requirements in robot soccer imply aspects of many different areas – including mechanical engineering, electronics, visual perception, artificial intelligence, multi-agent systems, and control – the strength of many robot soccer participants is their broad domain knowledge instead of a very narrow but deeper knowledge that typically experts from classical research have.

The scientific impact of developments in robot soccer differs from team to team. Due to the application-oriented workflow most approaches in robot soccer are adaptations of existing work to the needs of robot soccer. However, the special application domain might also create new solutions which lead to new scientific insights. We exemplified this relationship in section 3.1 with the development of a consistent description of the environment and the setup of a new kind of stereo camera. However, in most cases a lot of work is necessary to lift developments in robot soccer to an abstract, theoretical level so that they can be published in journals and proceedings and that they become visible outside of the original community. Here, robot soccer works similar to cooperations between researchers and partners from industry. However, publication of developments from robots soccer are not restricted by non-disclosure agreements and delays due to patent application.

As mentioned before one shortcoming of robot soccer is its unsatisfactory presence in the scientific community. New ideas are often not dissemanated or they are visible only in the robot soccer community. Often, publications about robot soccer do not focus on new ideas and lift them to an abstract level but just describe robot soccer in general or the overall design of a robot soccer team.

Certainly, the problem of unsatisfactory publications might also be related to the fact that the teams have to spend a lot of effort in making their teams competitive and maintaining them on a competitive level. New approaches can be implemented only after building a reliable software and hardware framework which requires at least two years of hard work. Moreover, since many participants are bachelor or master students they can only participate up to three years in a team until they finish their degree. Therefore it is difficult for many teams to organize a constant further development over time and to reach a level at which new ideas can be implemented.

4.2 Ideas for Future Development

To overcome the beforementioned difficulties there is the need to increase the ratio between scientific results and the effort in terms of manpower and money spent. One way could be to make achievements better visible to the scientific world outside of robot soccer. E.g. most websites of robots soccer teams show nice pictures and videos about their participation but only a few websites describe the individual research goals of the teams and the scientific achievements. Neither the scientific world nor the general public have the chance to get some ideas of what is the level of development in each league. Even for robot soccer participants the further development in other leagues remain vague. Furthermore, the results from the technical challenges are not available online, even some tournament results are missing. Here, a better presentation would already help a lot. It would also be a good idea to present the best scientific improvements of the present year in a special issue of a journal that has a better visibility than the RoboCup symposium. Moreover, it would foster scientific work of the RoboCup teams. For comparison, the journal of field robotics has reserved three issues for the teams which reached the finals of the DARPA urban challenge.

A second point to improve the scientific impact of robot soccer is to reduce the effort that is necessary to obtain a competitive level, especially for new teams. Here, more cooperation is necessary to avoid parallel developments and to reduce the engineering costs. One way is to foster open source development and open source dissemination of construction drawings. E.g. in the RoboCup simulation league the winning team has to provide its binary to other teams. Moreover, some teams are also publishing their sources code. In other leagues like the MiddleSizeLeague this practice has not yet been established in common. However, some teams have started to provide their source code like the teams *Brainstormers Tribots* in 2005 and *Carpe Noctem* in 2006. Certainly, new developments will stay proprietary for some years. However, for techniques that are established a dissemination of source code is reasonable. E.g. techniques for visual self-localization on a soccer field where in the focus of research during the years 2002–2005. Meanwhile, a couple of successful implementations exist and dissemination of these implementations could help new teams as a foundation for their own developments.

A third point is concerning middle-term plans for the further-development of the robot soccer leagues. At the moment rule changes are decided from year to year and sometimes the teams are surprised about changes that are announced five or six month' before the tournament. Certainly, this practice is inefficient for the teams since they cannot make middle-term plans for their own further-development. It would be very helpful if an agreement could be established of which rule changes will be done until 2015 (e.g. field size, playing in a hall or outdoor, development of set-plays, and communication with the referee) and what is the intended state in 2020 so that the teams have a guideline for their own research activities.

Finally, a last point is the financial aspect of robot soccer. Not only the development and maintenance of the robots but also the travel and transportation costs of the teams to participate at the international tournaments are immense. E.g. the expenses of the team Brainstormers Tribots in 2007 to participate at the German Open tournament and the RoboCup world championships summed up to 30,000 Euros. Although no new robots were build the expenses for spare parts, batteries, etc. summed up to 10,000 Euros or more. Obviously, no research group can provide these sums year by year. Therefore, it is important to find a better balance between expenses and scientific benefit. A biannual world championship with interleaving biannual regional competitions (e.g. European championships, American championships, etc.) could be one possibility to reduce the costs and to give more time to the teams for scientific dissemination and larger progress per tournament.

To summarize, participating in competitive research events like RoboCup can foster research in the respective area and induce new ideas and solutions. However, it is important to take the time to lift the developments onto a theoretical level to achieve scientific impact and to present results to the scientific world outside the own community. A participation should be intended to run for several years to benefit from the initial effort to reach a competitive level and it should be related to the core research interest of the research group.

Acknowledgements

The first author wants to thank the German Research Foundation (DFG) for being supported within the transregional collaborative research center SFB/TR 28 Kognitive Automobile.

References

[Baker and Nayar, 1998] S. Baker and S. K. Nayar. A theory of catadioptric image formation. In *Procedings of the Sixth International Conference on Computer Vision*, pages 35–42, 1998.

[Gabel and Riedmiller, 2007] T. Gabel and M. Riedmiller. Adaptive reactive job-shop scheduling with learning agents. *International Journal of Information Technology and Intelligent Computing*, 2(4), 2007.

[Gabel *et al.*, 2006] T. Gabel, R. Hafner, S. Lange, M. Lauer, and M. Riedmiller. Bridging the gap: Learning in the robocup simulation and midsize league. In *Proceedings of the 7th Portuguese Conference on Automatic Control (Controlo 2006)*, 2006.

[Hafner, 2003] R. Hafner. Reinforcement Lernen an einem mobilen Roboter. Master's thesis, University of Karlsruhe, 2003.

[Kitano *et al.*, 1997] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.

[Lauer *et al.*, 2005a] M. Lauer, S. Lange, and M. Riedmiller. Calculating the perfect match: an efficient and accurate approach for robot self-localization. In *Robocup 2005: Robot Soccer World Cup IX*, pages 142–153, 2005.

[Lauer *et al.*, 2005b] M. Lauer, S. Lange, and M. Riedmiller. Modeling moving objects in a dynamically changing robot application. In *KI 2005: Advances in Artificial Intelligence*, pages 291–303, 2005.

[Lauer *et al.*, 2006] M. Lauer, S. Lange, and M. Riedmiller. Motion estimation of moving objects for autonomous mobile robots. *Künstliche Intelligenz*, 20(1):11–17, 2006.

[Lauer *et al.*, 2009] M. Lauer, S. Lange, and S. Welker. Motion estimation from monocular and binocular observations in a mixed omnidirectional stereo camera system. *Robots and Autonomous Systems (submitted)*, 2009.

[Lauer, 2006] M. Lauer. Ego-motion estimation and collision detection for omnidirectional robots. In *RoboCup 2006: Robot Soccer World Cup X*, 2006.

[Merke *et al.*, 2004] A. Merke, S. Welker, and M. Riedmiller. Line base robot localisation under natural light conditions. In *ECAI Workshop on Agents in Dynamic and Real-Time Environments*, 2004.

[Müller *et al.*, 2007] H. Müller, M. Lauer, R. Hafner, S. Lange, A. Merke, and M. Riedmiller. Making a robot learn to play soccer using reward and punishment. In *KI 2007 Advances in Artificial Intelligence*, pages 220–234, 2007.

[Pfeifer and Bongard, 2007] R. Pfeifer and J.C. Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT Press, 2007.

[Pin and Killough, 1994] F.G. Pin and S.M. Killough. A new family of omnidirectional and holonomic wheeled platforms for mobile robots. *IEEE Transactions on Robotics and Automation*, 10(4):480–489, 1994.

[Riedmiller *et al.*, 2007] M. Riedmiller, M. Montemerlo, and H. Dahlkamp. Learning to drive a real car in 20 minutes. In *Proceedings Frontiers in the Convergence of Bioscience and Information Technologies (FBIT)*, pages 645–650, 2007.

[Riedmiller *et al.*, 2008] M. Riedmiller, R. Hafner, S. Lange, and M. Lauer. Learning to dribble on a real robot by success and failure. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2207–2208, 2008.

[Thrun, 2002] S. Thrun. Particle filters in robotics. In *Proceedings of the 17th Annual Conference on Uncertainty in AI*, pages 511–518, 2002.

[Voigtländer *et al.*, 2007] A. Voigtländer, S. Lange, M. Lauer, and M. Riedmiller. Real-time 3D ball recognition using perspective and catadioptric cameras. In *Proceedings of the third European Conference on Mobile Robots*, 2007.