Robot Soccer in Education

Henrik Hautop Lund

LEGO Lab, InterMedia

Department of Computer Science

University of Aarhus

Aabogade 34, 8200 Aarhus N., Denmark

hhl@daimi.au.dk

http://www.daimi.au.dk/~hhl

Paper category: Survey for Advanced Robotics journal Special Issue on RoboCup

Abstract

The educational experiences from using robot soccer in computer science courses are outlined. The educational approach can be termed *guided constructionism* and it differs from the traditional (pure) constructionism, which can be termed *unguided constructionism,* by combining hands-on experience with lecturing and guidance. We believe that some essential arguments needs to be taught through lecturing and guidance, but also find it essential for the students to actively participate in our robot soccer competitions when the aim is to educate the students to be able to make real world applications when graduating from the university. In the computer science courses, the students were taught to work with robot soccer players with pre-defined robot morphology, with modifiable robot morphology, and with robot team play. This should allow the students to learn 1) to manage the non-deterministic characteristics of the real environment, 2) to integrate hardware and software solutions, and 3) to manage collective and distributed systems. Indeed, with this knowledge, the students were able to win a number of international robot soccer tournaments.

1. Introduction

The development of information technology in the forthcoming decades will demand informatics to leave applications that are exclusively computer-based and instead occupy itself with the development of daily life objects (e.g. white goods, cars, etc.) that contains "intelligence". A number of skills is needed in order to construct such objects. The skills integrate different competences and not only the ability to make a perfect software application. In general, skills are needed to construct hardware objects that are governed by an "intelligent" or "adaptive" software that continuously interacts with a changing and non-deterministic real world. Our approach consists in teaching these competences to students by a constructivistic approach (learning by doing) that demands the students to confront themselves both teoretically and practically with a whole series of aspects.
In order for students to learn these skills, we have used robot soccer in education for a couple of years at the Department of Computer Science at University of Aarhus. We use mobile robots beacuse they are systems that can be programmed (and later on evolved, and therefore we can apply recent advanced in artificial intelligence and artificial life) and they are completely situated in the real world. Therefore, they are good tools for the study of a hardware-software relationship. Further, we use the robot soccer task, because it demands programming/evolution of single objects (mobile robots) and the study of diffusion of information in collective and distributed systems (robot teams).
In the courses *Adaptive Robots* and *Realization of Robots*, the Masters students are given the task of building robot soccer players to play in internal and external tournaments of robot soccer. For example, our students won the FIRA World Championship in S-KheperaSot and The Second Autonomous Robot Soccer Tournament, and we presented our LEGO robot soccer game at RoboCup'98 in Paris. Further, we have arranged other robot competitions for high school pupils in collaboration with a number of high schools. We have chosen to use robot soccer in education because it is a dynamic and complex game that can be played on a number of different levels, and students have to solve a number of different problems in order to develop successful robot soccer players. Here, these problems and the way students tend to solve these problems will be described with the focus on the educational process that the students go through. Section 2 describes three different robot soccer set-ups that we have used. The three set-ups have been designed as a three step educational process in which we work with 1) robots with a pre-defined robot body plan, 2) robots with a modifiable robot body plan, and 3) complex robot behaviours (such as team play). The description in Section 2 is technical, while Section 3 describes the educational aspects that we have experienced when using robot soccer in education. Section 4 summarizes related work and discusses how it relates to this work. Especially, the constructionism approach is confronted with our approach. We define the traditional constructionism approach as developed by Papert [1, 2] and, for instance, used by Martin [3] as *unguided constructionism*. Our approach can be termed *guided constructionism*, because it suggests to combine the constructionism approach with other methods (guidance) in order to allow the students to acquired knowledge in the most profound way.

2. Robot Soccer Set-up

We have worked with a number of different robot soccer set-ups. All these set-ups have an educational purpose, and are used in different courses. The three set-ups were designed as a three step educational process in which we work with 1) robots with a pre-defined robot body plan, 2) robots with a modifiable robot body plan, and 3) complex robot behaviours (such as team play). By going through this process, the students first get a thorough understanding of robot programming, secondly an understanding of the body and brain relationship, and thirdly an understanding of communication and distributed systems. In general, the students should learn 1) to manage the non-determinstic characteristics of the real environment, 2) to integrate hardware and software solutions, and 3) to manage collective and distributed systems.

2.1 Learning to manage the non-deterministic characteristics of the real environment (Pre-defined Robot Body Plan).

Every year in December, we arrange the Danish Championship in Robot Soccer. The finals always attract a big audience, who are entertained by the robot soccer play which can be seen also on a big screen and is broadcasted live on the Internet via mbone. Normally, we have around 9-12 teams participating in the Danish Championship in Robot Soccer. Each team consists of 2-3 students, who are in their 4^{th} or 5^{th} year of the computer science study. Therefore, they have a strong computer science background, but often no or little experience with real world systems. One of the main purposes with our robot soccer competitions is to provide the students with knowledge about real world systems, and our experience tells us that it is very difficult to achieve this knowledge without hands-on experience.

Figure 1. The single Khepera robot that plays the one-on-one robot soccer. © H. H. Lund, 1998.

The Danish Championship in Robot Soccer is a one-on-one competition, i.e. there is one robot on each team. By having only one robot in each team, it is obviously not possible to develop any team play, as is the case in the RoboCup and FIRA competitions. The reason for a one-on-one game is, that we believe the students should focus on the real world interaction (e.g. sensor interpretation, motor characteristics, environmental noise) and achieve a profound knowledge about this, before trying to develop more complex robot behaviours, such as team play. If one allows more complex behaviours, there is the risk that the students use all their time in developing team play. Such a team play might work well in simulation, but will probably fail dramatically in the real world, because the students have not used the (long) time it takes to understand the robot characteristics in the real world.

Hence, we agree with Brooks [4] that experiments must be performed in the real world. Our experience with the robot soccer competitions tells us that if we allow complex behaviours such as team play, then the students will tend to use a long time on abstractions and simulations. Therefore, they will fall back to a classical AI approach of abstraction and symbolic representation, and a control structure that in some sense resembles the classical AI sense-represent-plan-act control structure. When finally implemented on the robots, these approaches tend to result in robots that loose their robot soccer games. This is because they do not incorporate the environmental noise in the (symbolic) representation, and further, because of the long sense-represent-plan-act cycle, the robots are so slow that the opponent robot will have pushed the ball away before such a robot reaches the ball.

Each team has to develop a robot soccer player with the Khepera robot. The Khepera robot is a miniature mobile robot with 8 infra-red and ambient light sensors and two independent motors. It can be programmed in C with a protocol, and the cross-compiled program can be downloaded to run onboard the Motorola 68331. The Khepera robot has a modular structure, so extra devices can be attached on top of the robot. For the robot soccer competitions, we add the K213 vision module. It is a very simple 1D vision module that consists of 64 photoreceptors that each provides 256 grey-levels. The vision module has a view-angle of 36 degrees in front of the robot. Hence, the robot's vision with only 8 infra-red and ambient light sensors, and 1-dimensional 36 degrees vision is very simple, and the robot soccer behaviour will be limited by this. For example, the robot cannot see the goals, the opponent and the ball at all times, since the vision provides only 36 degrees view-angle.

The characteristics of the sensors defines the characteristics of the environment, since the, in some sense, "poor" vision does not allow distinction between all colours. Through empirical tests with the Khepera robot's vision module, we found the following to be a suitable environment. The ball to play with is a white or yellow tennis ball. The tennis ball is big enough to be seen by the vision module on top of the robot, and it is light enough to be pushed by the miniature robot. The walls of the arena are painted with a grey colour, and the goals are painted with a black colour. With these colours, it is possible to distinguish the different objects fairly well by simple vision processes. But, of course, the perception of colours depends on the light conditions. We try to provide uniform light conditions in the arena by having projectors to project light to the ceiling, and it is this light that will be reflected down on the arena. However, the environmental conditions (e.g. light) will change, and we should not do too much to avoid this, because this is one of the characteristics about the real world, and the students should learn about this. Even though the students are told this several times during the lectures, they are always amazed by the changes in environmental conditions when they are building their robot soccer player.

The players can be dressed with white and black striped paper, so that a robot can detect an opponent robot. However, the students very seldom reach the point where they have to consider the movements of the opponent. Therefore, so far in the Danish Championship and in international competitions, we have played without this white and black striped paper around the robot, but will probably do so in future when the robot behaviours become more complex.

In the first Danish Championship in Robot Soccer, we experienced that many periods ended with the robot(s) pushing the ball into one of the four corners of the playing field. The robot(s) was unable to move the tennis ball out from the corner, partly because of the circular shape of the Khepera robot. Some students tried a number of different strategies to move the ball out from the corner, e.g. turning very fast around itself while having contact to the ball, but these strategies were only successful in some particular cases. Hence, the shape of the playing field was changed for the subsequent tournaments by cutting the corners. This allows a much better flow of the play.

Figure 2. Students occupied with a preparation match for the Danish Championship in Robot Soccer. © H. H. Lund, 1998.

The above shows that a lot of empirical tests are necessary in order to make an environment that is suitable for the educational purpose that one has in mind. Often, one has to manipulate the light, colours and shape of the playing field, based on many empirical tests. In our case, we were interested in teaching computer scientist students about real world application, and therefore our set-up was biased by the educational purpose. For instance, we chose to allow only the same kind of robot on each team. In other RoboCup and FIRA competitions, the different teams are allowed to use any kind of robot that they build or buy (within a specific size). When using one specific robot, we ensure that the focus becomes on the computer science aspect of the application, rather than on more engineering aspects of robotics. All teams have the same motor characteristics, the same vision, etc. to work with, and therefore we can directly compare the properties of different robot controllers and different developmental processes.

In a sense, in this game, we are trying to limit the number of free parameters by using the same hardware platform for all robot soccer players. This allows the students to better focus their work at this initial stage when they have to learn about real world applications for the first time. Later in the educational process, the students are allowed to work with LEGO Mindstorms robots, with which they can investigate also the construction of the robot body plan.

2.2 Learning to integrate hardware and software solutions (Manipulating Robot Body Plan).

In the semester after having worked with the robots with pre-defined robot body plan for the Danish Championship in Robot Soccer, the students are allowed to work with manipulating the robot body plans. With *robot body plan*, we intend the type and number of sensors, the placement of sensors, the shape to the robot, the motors of the robot, the wheels (or other) of the robot, etc. This study is important and the students should learn about this relationship, since the performance of a robot controller is critical dependent on the robot body plan. This can, for instance, be shown with an evolutionary approach when co-evolving robot controllers and robot body plans [5, 6].

For this educational purpose, we use the LEGO Mindstorms robots. We find this robot platform a good choice in this educational context, because it allows students with no engineering background to assemble robots with different body plans. The robots are easily assembled with the well-known LEGO system of clicking bricks together. For instance, the central control unit has three input channels and three output channels to which one can attach sensors and motors. The sensors and motors are cable connected with cables that have small LEGO bricks on the ends, and they are therefore straight forward to click onto the control unit. This can easily be done by pupils down to the age of 10.

With this tool, the students can investigate the relationship between body and brain. The controller ("brain") can be programmed with a graphical programming language by younger pupils, while students at higher levels can program the LEGO Mindstorms robots in a Windows-based programming language such as Visual Basic, Visual C++, Visual

Java++, etc. Further, for the advanced computer science student, it might be possible to write programs directly to the microchip.
In the *Realization of Robots* course, the students are again given the task of making robot soccer players, but this time they are allowed to develop both the controller and the morphology of the robot. However, they are confined to use the LEGO robot platform, and hence only very primitive sensors. The available sensors include light sensors, angle sensors, temperature sensors, and bump sensors. Of course one can build other sensors to add to the robot. For instance, we have made bend sensors, ears, etc., but for the autonomous LEGO robot soccer, we chose to allow only the other sensors, so that the study of the functionality of these would be more thorough.
The few different kinds of sensors certainly limits the possible behaviours, but again it makes the students focus more on the computational aspects, and, for example, makes the students study ways to combine sensor information.
In this case, the arena was approximately 2m*3m, and had different colours on the floor in different areas, so that one could make a robot navigate around in the arena according to the colours perceived with a light sensor. The ball was a white plastic football, which gave a good reflection when light was projected onto the arena.
Even though that we carefully manipulated the characteristics of the arena, it was difficult for the students to design robots that would make as good a soccer game as was the case with the Khepera robots, though some of them succeeded in making robots that could find the ball and push it into the opponent goal. But importantly, by going through the process of both building and programming the robot, the students obtained new knowledge on top of what was experienced with the robot soccer with pre-defined robots. For instance, the students experienced that the principles for development of controllers for pre-defined robots had to be modified in order to apply to robots with modifiable morphology.

A video of the autonomous LEGO robot soccer is available on the following web-site:
http://legolab.daimi.au.dk

2.3 Learning to manage collective and distributed systems (Team Play).

We wanted to show our work on robot soccer at RoboCup'98, which was held in Paris in the summer of 1998 during the World Cup in Soccer -- France'98. Even though the students were able to make fully autonomous LEGO robot soccer players, we did not find the results impressive enough to show to the audience at RoboCup'98. Given the right environment, some of the fully autonomous LEGO robot soccer players could find the ball and push it into the goal, but we did not find the results robust enough to ensure a lot of goals being scored in each match. Further, for a more complex behaviour such as team play, we needed other (vision) sensors and communication.
Therefore, we chose to work on another LEGO robot soccer set-up for RoboCup'98. In this set-up, we added an overhead vision system, and communicated the information obtained with this vision system to the LEGO robots via infra-red communication. Each team consisted of a goalkeeper and two robot field players. The overhead vision system was used because it allowed us to obtain information on position of players and ball in an

easy fashion. However, on the long terms, we believe that the RoboCup research goal should be achieved without such an overhead camera, because there are complex problems involved in active (onboard) vision that might be directly related to the development of intelligent robots. But, based on the students' tests with the fully autonomous LEGO robot soccer, we chose to use the overhead vision system for the 1998 demonstration, since it allowed us to develop more complex behaviours at this moment. Indeed, with this set-up, we achieved a good soccer game, and goals were scored in all matches.

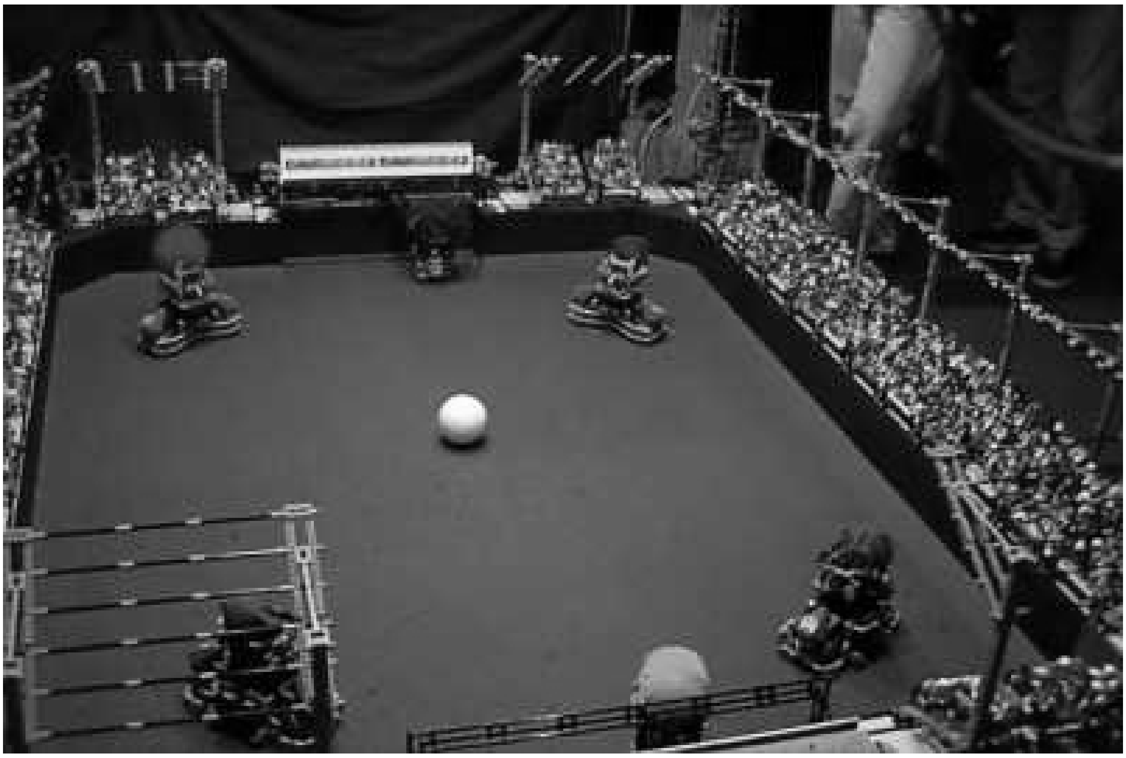

Figure 3. The LEGO robot soccer set-up that was used during RoboCup'98. There are two field players and a goalkeeper on each team. © H. H. Lund, 1998.

Further, to put the robot soccer play into an appropriate context, we also built a huge LEGO stadium on which the robots played. The stadium had a lot of functionality built into it. For instance, it included functional light towers, rolling commercials, score board, small scanning cameras that projected images up to big screens. It also included a sound system that could produce different sounds for shooting, tackling, whistling, cheering, goal-shouts, etc. Further, the stadium had almost 1500 small LEGO spectators who could perform the "Mexican wave". The wave was performed using the principle of emergent behaviour. When rising, a sensor underneath one section of LEGO spectators would trigger the next section of LEGO spectators to rise. Hence, each section had a local control based on the input from the previous section, and the global behaviour of the

“Mexican wave” emerged from the interaction between the sections with local control. This phenomena and the stadium set-up has been reported thoroughly in [7].

3. Educational Experience.

When starting on the robotics courses, most students have studied computer science for 3-5 years. Therefore, they come to the robotics courses with a thorough knowledge about computer science, both a theoretical knowledge and pratical programming knowledge from programming compilers and the like. However, the students’ knowledge about real world applications with control of devices external to the computer is, in most cases, none. Therefore, the students have a number of ungrounded expectations of the perfomance of such systems. Often, these expectations are based on the students’ previous experiences with programming in deterministic environments in the computer. Also, it partly arises from the whole natural science approach in which we have a profound belief that systems can be broken down into smaller systems that each has a determistic functionality. Therefore, the students must go through a number of lectures and practical experiences before they are convinced to change their unrealistic view of real world applications.

The students’ belief in a deterministic reality can be observed in their robot soccer projects, and we have experienced that the students change this view by going through the process of building robot soccer players. Some of these experiences can be documented by looking at the progress reports that the students are asked to hand in during their projects.

First of all, when starting the educational process, the students have totally unrealistic beliefs of the capabilities of the robots. For the one-on-one robot soccer project with the Khepera robots, one of the groups of students wrote[1] in their first plan for the robot development:

> “... Next phase is to locate the ball (by turning from side to side), and rush towards it, in the hope of/attempt to hit it. Then, rush back again and “start all over”. Here, it might be advantageous to decide the distance to the ball, and only drive after it when it is “close enough”, dependent on how well this can be decided from the response of the camera (the instrumentation).”

Obviously, these students had no idea about the capabilities of the simple 1-dimensional grey-level vision module of the Khepera robot. Their general idea was that they would be able to have the robot to play some kind of airhockey, but a couple of weeks later, after making experiments with the robot and the vision module, they changed view:

> “Our original strategy of allowing the robot to play “airhockey” has proven totally impossible.”

In general, most students go through the process of having to change their initial, general strategy when they achieve more experience with working with the real robot in the real

[1] All quotes from student reports have been translated from Danish to English by the author.

environment. Another example is, that almost all students believe that they can incorporate a very precise global positioning system. For instance, one group writes:

> "... Further, we want to have a "sense of locality" that tells something about the angle that the robot has in relationship with the opponent goal. We want to implement this by using a counter that keeps track of how much the wheels have moved."

Apparently, at this stage, these students (and most other students) have no idea about the role of friction, spinning, etc. But, nonetheless of these problems that they should run into, they write a couple of weeks later:

> "It looks like we will get the sense of locality to work. But we have observed that when the ball hits the robot, the robot will make a pirouette."

But after failing miserably in the robot soccer tournament, these students write:

> "We never reached the point of having a global direction to allow the robot to know where it was supposed to score."

However, it is alarming that they have not realized the profound problems with such a global position system, though most other groups realize these problems, and come to similar conclusions as this other group:

> "We have realized that very little interference will make us loose the orientation. It will quickly fail, if another robot drives into us, or if we stand still and push against the ball or the wall."

Therefore, most groups realize the difficulties in making a global positioning system work in reality, and start to consider how to use environmental feedback to update an approximate positioning system. The students are used to have all information available in a pure form in the simulation work that they made previously in the computer science study, but now they experience that this information will by no means be available in a pure form in real world applications. Therefore, they change their view, and start to think about how to make use of the little knowledge that they might obtain via feedback from the environment. But they also experience difficulties when trying to obtain feedback from the environment, because the sensors are almost always much more primitive than the students expect. Most groups write similarly to this group:

> "First of all, we were surprised by how much importance the robot's surrounding has. We expected that as long as one put a camera on top of the robot, it would be no problem to allow the robot to see."

As Martin [3], we notice that, initially, also our students believe that sensors give a clear and unambiguous input, and that they therefore can use abstraction. The abstraction and pseducode is used instead of experimentation in order to overcome software complexity, and only later through the experimentation, the students realize the true nature of the sensors. In a sense, this resembles the discussion about the classical AI approach to robotics, in which the hardware and software tasks were believed to be distinguishable, so that engineers could work on the hardware, while AI researcher could make abstractions and work on the software only in order to create an intelligent system (robot). Nowadays, a number of new AI approaches to robotics reject this division and focus on embodied AI. Through the experimentation, the students realize that their abstractions do not hold:

> "In the semifinals, our robot had difficulties in recognizing the ball. It resulted in the robot moving backwards towards to the end line of the field and just stood still while staring. It is not inconceivable that the light conditions were changed just slightly, but enough for the thresholds values from the training room not to be true anymore. It is notable, that our robot performed better when the ceiling light was switched off."
> [...]
> "The robot should also, in an adaptive manner, decide the threshold values for object recognition, so other than the usual and known light will not hamper as much as was the case during the semifinals."

A lot of students experience exactly the same, because they are used to making abstractions and believing that the sensor information can be interpreted in a straight forward manner using fixed thresholds to classify the input. They experience that the approach works one day, but fails miserably the next day when, for example, it is no longer sunny weather outside the lab window. Initially, the students become frustrated and blame the hardware. But the long process of experimentation makes the students realize that the approaches that have been taught in the Adaptive Robots lectures can be used to overcome this problem. In fact, the winning robot from our competitions had adaptive thresholds that would change during the play. With this robot, our students went on to win both the FIRA S-KheperaSot World Championship 1998, and the Second Autonomous Robot Soccer Competition in 1998, probably with the only robot (also in most RoboCup teams) that actually changes behaviour in an adaptive manner during the play. We believe that the success of our robots in the international competitions is largely due the fact, that our students have built up a profound knowledge about adaptive robot behaviour and are able to use this in their robot soccer players.

During the process, the students also learn that the battery charge might influence the behaviour of the robot, and therefore influence the abstractions that they used initially. The students are always encouraged to make numerous tests without cable connection whenever possible, but nevertheless most groups of students are surprised by the differences, as for instance this group:

> "We had been "training" extensively with our robot and Tuesday night we were quite satisfied with his behaviour. He could, consistently recognize ball, goal, and wall, he could find the ball, and then he could dribble the ball in goal if he got the chance. It was all very edifying. We chose to train with the cable

> connected partly to avoid problems with powering, partly to observe our debug messages. It turned out not to be a smart idea, because when we arrived at the preliminary rounds in which we ran without cable connection, suddenly there was something that did not work as usual. In our first match, the vision was reduced, probably because of low battery level."

Again, the problems arise because the students believe in being able to make abstractions and, for example, classify sensory data with fixed, pre-defined thresholds. However, the sensory data will vary depending on a lot of internal and external conditions that the students can only learn about through hands-on experimentation (combined with theoretical lecturing).

Therefore, our educational experience with the robot soccer competitions tells us that we must combine lecturing with hands-on experiments. This is because students will have difficulties in believing in another view on the real world than the traditional view of a world with deterministic characteristics in which abstraction is feasible. It seems like the students are only able to change this view when they are actually experiencing themselves, a lot of times, that their robots will fail with a control program that depends on the abstraction. Then the students start changing their view and implement the more adaptive solutions that they have been taught about in the lectures.

4. Related Work and Discussion

At the Department of AI at the University of Edinburgh, the mobile robot group has annual LEGO robot competitions. This is part of the Intelligent and Sensing Control course by lecturer John Hallam. In this course, the students' final project is to participate in a competition like robot sumo wrestling or robot rugby. The aim of the competition is to give students hands-on experience with robot building and programming. This, in combination with the lectures, should allow the students to acquire knowledge about relationship between sensing, acting, control, etc. and how all these competences might influence intelligence. This knowledge is viewed as essential to the students' artificial intelligence study. We hope that our approach at the University of Aarhus with robot soccer competitions will allow our students to acquire a similar knowledge.

There exists a number of open robot competitions such as Micro Mouse and FIRST competitions. Micro Mouse has been running since the late 1970s, and it consists of designing an autonomous robot known as a "mouse" which should navigate its way through a maze. The mouse has no prior knowledge of the maze configuration before its release in the maze. During three runs, the robot should travel from the starting point to the center of the maze. The first two runs are used for data gathering and the final for a high speed run, to obtain the fastest handicapped time.

FIRST (For Inspiration and Recognition of Science and Technology) is a non-profit organization whose mission is to generate an interest in science and engineering among youth, and their primary means of accomplishing this goal is through annual robot competitions, which began in 1992. The final of these competitions are held in the Walt Disney World Epcot Center. The competition is an engineering contest in which high

school students team up with engineers from businesses and universities. In six intense weeks, students and engineers work together to brainstorm, design, construct and test their robot. The aim of the FIRST competitions is to bring together schools, businesses, and universities, and thereby provide an exchange of resources and talent, highlighting mutual needs, building cooperation, and exposing students to new career choices. In the process, all participants should discover the important connection between classroom lessons and real world applications. For instance, in the autumn 1998, the LEGO Lab arranged a FIRST LEGO League Pilot for Danish school classes with pupils at the age 12-14. In this pilot project, the pupils worked together with tutors from university to develop their own robots to compete in a friendly competition against robots developed by pupils from other schools.
More information about a variety of different robot competitions can be found at:
http://www.ncc.com/misc/rcfaq.html

It is striking how much our educational experience (e.g. see quotes in the above section) resembles that of Fred Martin [3]. At MIT, Martin is also running annual robot competitions with the aim of giving students experience with robot building. The competitions are part of a hands-on, workshop-like course for undergraduates, and it runs (intensively) for one month only.
In fact, a number of our educational experiences with the robot soccer competitions reported in Section 3 have direct parallels to what Martin has experienced with his robot competitions at MIT. For example, Martin also suggests that his students run into problems because they tend to build robots that perform properly only under ideal conditions. He states clearly that “Students repeatedly build robots that are not well-equiped to deal with the exigencies of the real world, but rather with the specifications of an idealized, abstracted world -- a world that the robot designers would like to believe is a close representation of reality, but is not.” [3], pp. 189.
The approach taken by Martin is based on the theoretical considerations put forward by Seymour Papert [1,2]. These considerations are to a large extend natural extensions of the work by Jean Piaget [8], and they have been termed constructionism. Constructionism suggests that learning is achieved most effectively by participation in the construction of artifacts. The artifact becomes an “object to think with”, which can be used to explore and express ideas which are related to the field (the thing) under investigation. In the robot competitions, the students are allowed to construct their own robots, and learn a lot about real world applications by going through the building process.
However, we choose to have robot soccer competitions at different levels, because we believe the educational process to be slightly more complex than what is suggested in constructionism. In its most pure form (which we term *unguided constructionism*), constructionism seems to suggest that children should just be allowed to play with what they find fun, and they will learn what they need to learn by this play. We believe that, for some subjects, children/pupils/students need guidance, and that there exists subjects that are profound for a scientific knowledge which do not lend themselves to a constructionism approach, but will have to be taught in a more traditional manner. This is especially the case when educating graduate students. Constructionism was initially introduced in educating children (e.g. younger pupils), where it, in my view, also seems more

appropriate. But for education at higher levels, we need to combine constructionism with other pedagogical methods to ensure that the students obtain a profound knowledge about the subject under study.

Therefore, we here investigated *guided constructionism* and the use of a three step process with 1) pre-defined robots, 2) robots with modifiable robot body plan, 3) complex robot behaviour. Constructionism is the core of 2) and enters to some extend in 3), while it plays only a minor role in 1). With the three step process we (try to) ensure that the students firstly gets knowledge about programming in the real world and a tool to compare their different approaches. The comparison is facilitated a lot for the students by not allowing them to modify the robot body plan. The students obtain an essential knowledge about robot programming by working with the robot soccer project with pre-defined robots (combined with lectures). Having achieved this essential knowledge, the students are allowed to manipulate the robot body plan by "playing" with LEGO robots, and they obtain a knowledge about the relationship between controller and robot body plan (between brain and body). This second part of the three step process can be viewed as constructionism, but it is based on the knowledge that the students have achieved by previously going through the process of learning about robot programming with the pre-defined robots. In the third part, the basic knowledge about robot building and programming is used to allow a full real world application to be developed. But again, the constructionism is accompanied with study of different subjects (e.g. communication) in a more classical manner.

Hence, we believe that the constructionism approach should be combined with other methods, and that there exists essential arguments that are better acquired by the students with other approaches. However, we agree with the constructionism approach that hands-on experience must play a major role in education and that it often facilitates the students' acquisition of knowledge about an artifact, because the students will have experience with working with the artifact. Certainly, we find it essential for the students to actively participate in our robot soccer competitions when the aim is to educate the students to be able to make real world applications when graduating from the university.

5. Conclusion and Future work

In this paper, guided constructionism has been outlined and a number of empirical experiences using guided constructionism has been described. The test case has been robot soccer, which has worked as a three step process in which students 1) learn to manage the non-deterministic characteristics of the real environment, 2) learn to integrate hardware and software solutions, and 3) learn to manage collective and distributed systems.

Further, we believe that robots can be useful in education of a number of other subjects. For instance, robots can be used as an educational tool for artificial life and biological investigations, as described in [9]. Also in this context, the robot soccer competitions might be a test platform, since we can, for instance, imagine studying the evolution of robot controllers, the evolution of communication, the evolution of suitable bodies, etc.

Acknowledgements

Ole Caprani was co-lecturer in some of the robotic courses. Thanks to the students on the robotic courses who provided the reports cited in this paper. Orazio Miglino provided a number of highly valuable comments and inspiration. Thanks to John Hallam for inspiration regarding robotics courses.

References

[1] Papert, S. *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books, New York, 1980.

[2] Papert, S. Constructionism: A New Opportunity for Elementary Science Education. *A proposal to the National Science Foundation*, 1986.

[3] Martin, F. “Ideal and Real Systems: A Study of Notions of Control in Undergraduates Who Design Robots”. In Y. Kafai and M. Resnick (Eds.), *Constructionism in Practice: Rethinking the Roles of Technology in Learning*, MIT Press, MA, 1994.

[4] Brooks, R. A. “Elephants Don’t Play Chess” In P. Maes (ed.) *Designing Autonomous Agents*, MIT Press, MA, 1990.

[5] Lund, H. H., Hallam, J., and Lee, W.-P. “Evolving Robot Morphology” In *Proceedings of IEEE* 4th *International Conference on Evolutionary Computation*. IEEE Press, NJ, 1997.

[6] Lee, W.-P., Hallam, J., and Lund, H. H. “A Hybrid GP/GA Approach to Co-evolving Controllers and Robot Bodies to Achieve Fitness-Specified Tasks” In *Proceedings of IEEE 3rd International Conference on Evolutionary Computation*. IEEE Press, NJ, 1996.

[7] Lund, H. H., Arendt, J., Fredslund, J., and Pagliarini, L. “Ola: What Goes Up Must Fall Down” To appear in *Proceedings of Artificial Life and Robotics 1999 (AROB’99).*

[8] Piaget, J. and Inhelder, B. *La psychologie de L'enfant*. Paris, P.U.F, 1966.

[9] Miglino, O., Lund, H. H., Cardaci, M. “Robotics as an Educational Tool” To appear in *Journal of Interactive Learning Research*, 1998.