

KiRo – An autonomous table soccer player

Thilo Weigel, Bernhard Nebel

Institut für Informatik

Universität Freiburg

79110 Freiburg, Germany

weigel,nebel@informatik.uni-freiburg.de

Abstract. This paper presents KiRo – a system capable of playing table soccer on a competitive level and in a fully autonomous way. It can serve a human both as a teammate and an opponent but also allows for matches between two artificial players. KiRo introduces the table soccer game as a new domain for the research in the fields of robotics and artificial intelligence.

1 Introduction

The vision of the RoboCup research community is to build a robotic soccer team capable of beating the human world champion – by 2050 [3]. However, the games of the past RoboCup competitions show, that there is a long way to go before we even can think about playing against a human team. It would be clearly helpful if one could define tasks related to RoboCup that allow for a competition against humans without being forced to implement everything necessary to play real soccer. For example, one could define a mixed human/machine game for the simulation league. In fact, there have been such games already.

Even more interesting would be a game that is already well known and played by a lot of people. Such a game is *table soccer*.¹ This game is a popular pastime in bars and amusement arcades similar to billiards, but also a sport in its own right with competitions and even world championships. Due to the rather simple environment and the restricted skills required for playing it, it seemed realistic to us to develop an autonomous table soccer player capable of beating ordinary human players and to face the challenge of beating the human world champion before the year 2050.

Based on our experience of developing the *CS Freiburg* robotic soccer team [2,4–6], we started to design and construct an autonomous table soccer player, which we called *KiRo*.² The table soccer game raises research questions similar to the ones addressed in the RoboCup context. However, as the sensor and motor problems are easier to solve, table soccer allows one to focus on high-level research issues as in the simulation league. Nevertheless, it presents the challenge of constructing an integrated system and has the fascination of a real robot league. Among other things, it seems feasible to successfully apply learning techniques in a real world scenario and it seems worthwhile to explore how the behavior of the opponent can be modeled in order to refine the player’s strategy.

An autonomous table soccer player has the potential of fascinating people watching or playing against it. It joins entertainment and research and thus may be attractive for various purposes. Maybe there will be a “table soccer league” in the future?

¹ Table soccer is also commonly known as “Bar Football” or “Foosball”.

² *KiRo* stands for *Kicker-Roboter*, which is the german expression for “table soccer robot”

The rest of the paper is structured as follows. We start with a description of the general system architecture in section 2. Section 3 presents details of the hardware we developed to control the rotational and translational movements of a player's rod. Section 4 describes how we gain a model of the world from the camera data. Section 5 presents KiRo's basic actions and how they are selected. Section 6 reports results from matches KiRo played in the public. In section 7 we conclude and give an outlook on our future work.

2 General Architecture

An important criterion in the design of KiRo was to modify the natural environment only as much as absolutely necessary. Therefore we bought a commercially available table soccer and attached the units for controlling the rods to the outside of the table.




Fig. 1. The hardware setup

Each unit is individually mountable and controls one rod of a player. The only sensors used are absolute position encoders of the motors and a color camera overlooking the table³. A standard personal computer connects to the camera and control unit and handles all the information processing during a game. Figure 1 gives a sketch of the setup.

Figure 2 shows the structure of the control software we developed. From the camera data the position of the ball and the positions of the players are estimated. These estimates are integrated into a world model where the direction of motion and the velocity of an object are calculated. In addition, the world model generates information which helps to locate the ball if not visible. The encoders are used to detect if the ball is locked between a figure and the

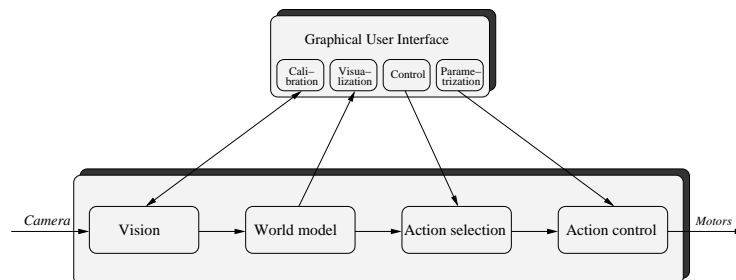


Fig. 2. The software architecture

floor. This information is very helpful if a player wants to stop and control the ball, e.g. for a pass to a figure of the same rod.

³ We discarded the idea of mounting a grid of photo sensors just above the playing surface for ball recognition, because this would have meant a considerable modification of the soccer table

Based on the world model the action selection modules decides which behavior from a set of basic skills is the most appropriate one for the given situation. The execution of the selected action is then monitored by the action control module which sends the suitable commands to the motors.

Table soccer can be a very fast game where the ball may move at velocities of up to 10 m/s. This means that even at a frame rate of 50 Hz the ball possibly travels as far as 200 mm between two consecutive frames. In order to cope with these demanding conditions we attempt to minimize the delay between a change in the environment and an action being taken in response. By synchronizing the vision and action selection processes we achieve that an action is selected and executed as soon as a new camera image is available. This reduces the delay to the amount of time needed for digitizing the camera's analogue signal (approximately 20ms at a very high shutter speed) plus the time needed for processing the new information and communicating with the motor controllers (currently approximately 11 ms). In order to cope with this delay, we date back by 20ms the timestamp which is assigned to a newly arrived image and estimate the positions of all objects 11 ms ahead of the current time. Of course, this increases the uncertainty of an object's position estimate, but decreases the effects of the delay.

But reactivity isn't everything anyway: Interestingly, professional players affirmed, that very often their strategy envisions just *not* to react to the movements or posture of the ball or the opponent figures. At an advanced level of play one would be fooled by an opponent's trick too easily and it is therefore advisable to concentrate on an elaborate and constantly changing positioning scheme. What distinguishes a professional player is rather his strategical play and his motoric skills to manipulate the ball than his fast reaction time.

3 Hardware

Because in Germany *Löwen* soccer tables are commonly used both for competitions and recreation we decided to use their *Home Soccer* model⁴ for our purposes. The table provides a playing field of the size 1200mm x 680mm and has four rods on each side. There are figures attached to each rod (1 for the goal keeper, 2 for the defender, 5 for the mid field and 3 for the attacker), which can be used to manipulate the ball.




Fig. 3. A control unit

We designed the units for controlling the rotational and translational movements of a player's rod in such a way that they are easily mountable to other manufacturer's tables as well. Figure 3 shows a model of a control unit with a player's rod attached to it. Coupled to a belt drive, the slide glides on a spline shaft when the belt is moved by a motor. A second motor and belt drive cause the spline shaft, and thus the belt pulleys of the slide, to rotate. As the player's rod is attached to

⁴ <http://www.loewen.de/produkte/soccer/homesocc.htm>

the upper belt pulley of the slide, any desired translational and rotational movement of a player can be achieved by controlling the corresponding motors appropriately.

The *Faulhaber* servo motors and motor controllers we use are addressable via an RS232 link in a comfortable way. To control eight motors (2 for each unit) at the same time we equipped the PC *AMD Athlon, 700 MHz, 256 MB Ram* with a *Stallion Easy I/O* multi-port serial adapter.

Even though more specialized sensors could be used, we preferred to rely mainly on the position encoders of the motors and on one color camera. The vision system we employ consists of a *Phytec VCAM-110* camera and a *Phytec pciGRABBER-4* frame grabber, providing PAL images in YUV format. Because we analyze each half frame individually, we achieve an effective frame rate of 50 Hz with an image resolution of 384x288 pixels. In addition, a safety light grid is used to guarantee that nobody will be harmed when putting a hand between the player rods.

4 Vision and World Modeling

Using the very efficient CMVision library [1] for color segmentation we are able to extract the positions of the players and the position of the ball at a frame rate of 50 Hz. Since the position of the table relative to the camera may change during a game, we have developed a simple but effective calibrating mechanism which allows us to re-calibrate the system every few seconds. The transformation matrix from image coordinates to real world coordinates is determined by examining the field lines and finding both the center circle and the center lines using a simple template matching method. First, starting with default values, the position and the size of a circle template is varied systematically until the correlation with the camera image reaches a maxima. As in the successive cycle the




Fig. 4. An image from the overhead camera (a) and the world model derived from it (b). In the image the calibration circle and line are plotted black. In the world model the ball is marked as a white circle and the ball owner is marked by a white frame.

previously determined circle center, circle position and line orientation are taken as starting values, the described method usually finds the relevant field lines within a few milliseconds – if the soccer table doesn't move considerably. The circle's center then

yields the position of the table, the angle of the straight lines its orientation and the circle radius the zoom factor. Figure 4 (a) shows an image from the overhead camera. It can be seen that the calibration routine recognized the center circle and line.

A player's rod has only two degrees of freedom and both its position and the distance between its figures are known. Thus, individually detecting the positions of all the figures of a rod creates redundant information. This means that we can improve the estimates of the figures' inclination and position by averaging over all figures of that rod. The *a priori* knowledge also permits us to improve the estimates for an individual player because different blobs belonging to the same figure can be associated reliably.

The angle of a figure is estimated by comparing the width of its bounding box with the minimum and maximum possible values. Due to the limited resolution of the camera image only a rough and ambiguous estimate is possible. For instance, one cannot really tell if a figure is pointing right towards the table or in the opposite direction. However, by examining to which side of a rod a figure's bounding box further extends, it can usually be detected towards which side the figure is inclined. Knowing towards which side of a rod a figure is "up" and assuming that a figure is "down", when the bounding box and thus the figure's angle is small, is sufficient for our purposes at the moment.

Currently we use a yellow ball in order to simplify the perception problem. Since there is no other yellow object on the field, all yellow blobs can be used to calculate an overall bounding box. Usually its center corresponds to the center of the ball on the field, but there might be considerable inaccuracies if the ball is partially covered by a rod's figure. However, knowing the expected size of the ball, the bounding box can be reconstructed taking into consideration the covering figure.

Quite often the ball cannot be identified in the camera image. This happens either because a goal was scored or - more frequently - because the ball is obscured completely by a figure. In order to ensure an appropriate behavior even in these situations, we mark a figure as the "ball owner" if the visible ball is close enough to it. The knowledge, which figure was the ball owner before the camera lost track of the ball permits us to assume that the ball is located right underneath this figure, even if this figure is moving.

Figure 4 (b) shows a screenshot of the world model constructed from the camera image. Note, that even though not plotted in the world model, it is known, towards which side of a rod a figure is inclined.

5 Action Selection and Action Control

In the first prototype version of KiRo we employ a simple decision-tree like action selection mechanism. It considers the current game situation and chooses one of the following basic actions for execution:

- *DefaultAction* move and turn rod to default home position
- *KickBall*: rotate the rod by 90° to kick the ball forward
- *BlockBall*: move the rod so that a figure intercepts the ball
- *ClearBall*: move to the same position as *BlockBall* but turn the rod in order to let the ball pass from behind
- *BlockAtPos*: move the rod so that a figure prevents the ball passing at a specific position

- *ClearAtPos*: move the rod to a specific position and turn the rod in order to let the ball pass from behind

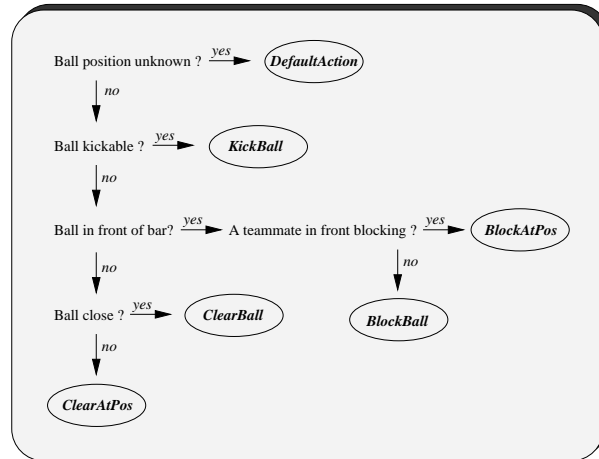


Fig. 5. The decision tree for selecting an action

Figure 5 shows the decision tree used to select an action for the rod of a player. Only a few predicates are used and the actions are selected in a straightforward way. All the rods are treated in the same way, except that the positions for *BlockAtPos* and *BlockBall* are calculated individually.

Every figure of a player’s rod can move within a certain range. Usually the ranges of two figures overlap and consequently there may be situations with two figures as candidate for moving to a specific target position. In these cases we select the figure with the best trade-off between its distance to the target position and the distance from its range limits to the target position. To avoid oscillation, e.g. when blocking the moving ball, we additionally prefer the figure that was selected in the last cycle already.

The exact target position where a figure tries to block the ball depends on the position, direction and velocity of the ball. If the ball is moving very fast, the position where the ball is anticipated to pass the rod is taken. However, if the ball is moving very slowly, the target position is set in the way of the expected kicking direction of the opponent: it is assumed that the figures of an offensive rod always aim at the goal and that other figures usually try to kick the ball just forward. In order to achieve smooth transitions between the anticipated and the expected position, the target positions for intermediate ball velocities are calculated by interpolating between the two positions.

To avoid unnecessary “hectic” behavior the velocity at which a rod is moved is determined by considering how urgent it is to reach a given target position. If the ball is far away and it is moving only slowly there is no need for a fast reaction. But the closer and the faster the ball gets, the higher is the speed assigned to the rods.

The currently available basic actions and the manner of selecting them lead to a very “agile” and effective behavior. The well organized team line up and the reliable

ball blocking skill make it quite difficult for an opponent to bring the ball forward. In turn, the fact that the ball is kicked forward as soon as possible makes it hard for opponents to react fast enough to block it. Test games have shown, that even with a simple strategy like the one presented here, beginners and amateurs can be beaten by the autonomous table soccer player.

6 Results

At the RoboCup German Open '02 in Paderborn and the Hannover Industrial Fair '02 KiRo played during nine days for a total of about 22 hours against human players. Mostly men, but also a lot of women were queuing up, curious to play table soccer player against a machine. In the same way kids, adolescents and adults were fascinated by KiRo, ambitious to win and obviously having fun when playing. We were surprised by the fact, that most of the "test players" had considerable experience and were already playing at a high standard.

Altogether we recorded 102 matches KiRo played against human opponents, ignoring a lot of "trial games", where people gave up very early or just wanted a quick impression of how it is like to play against a machine. Always two players were playing as a team against KiRo. For all the games we recorded the team's playing standard, which we assessed by observing the players' skills and their strategical behavior, distinguishing between the following categories:

- A *beginner* has hardly ever played and has neither special skills nor a strategical understanding for the game.
- An *amateur* plays once in a while, has fundamental skills and a strategical understanding of the game.
- An *advanced player* is playing regularly and consciously improving his skills and strategy while playing
- A *professional* is explicitly practicing special shots or tricks, plays with an elaborate strategy and competes at tournaments

	Games won : Games lost	Goals shot : Goals received
Beginners	17 : 0	154 : 28
Amateurs	56 : 11	501 : 284
Advanced Players	3 : 11	75 : 117
Professionals	0 : 4	7 : 40
Total	76 : 26	737 : 469

Table 1. game results

Table 1 shows the results of the games. Of course, only a rough and subjective assessment of the teams' playing standard was possible. KiRo won about 75% of all the games and proofed to be a challenge for beginners and amateurs. Even against two professional players (one of them the runner up at the European Championships in 1998), KiRo managed to score some lucky goals until they got used to its playing style. Nevertheless, like the audience and all the other players they appreciated KiRo and offered to contribute their experience and knowledge for improving KiRo's strategy.

7 Conclusion and Future Work

We presented a system capable of playing table soccer in a fully autonomous way. In its first prototype version it plays with a simple but effective strategy and is able to win against beginners and amateurs. People playing against KiRo were obviously having fun and with respect to the entertaining aspect (not the social aspect, though) of the game they saw little difference to playing against a human player.

The first positive feedback from human test players allows us to be optimistic regarding the acceptance of KiRo to the wider public. But the existing prototype is more than just a toy: it also represents a robust basis for research in the fields of robotics and artificial intelligence. In competitions against human players or other artificial players research can be evaluated in an attractive way in the real world.

Our future work includes improving the hardware and extending the vision system to enable playing with a white ball. In addition, we plan to focus on designing a new action selection mechanism which incorporates new skills, allows for different playing styles and adapts itself to the characteristics and capabilities of the opponent player. For tuning the parameters of the basic skills we attempt to apply learning techniques in the real environment.

Acknowledgments

This work has been partially supported by *Deutsche Forschungsgemeinschaft* (DFG), by *SICK AG*, and by *Faulhaber GMBH & CO.KG*. The authors would also like to thank *Joachim Koschikowski* for the technical drawings of the control units.

References

1. James Bruce, Tucker Balch, and Manuela M. Veloso. Fast and inexpensive color image segmentation for interactive robots. In *Proc. of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '00)*, volume 3, pages 2061 – 2066, October 2000.
2. Jens-Steffen Gutmann, Wolfgang Hatzack, Immanuel Herrmann, Bernhard Nebel, Frank Rittinger, Augustinus Topor, Thilo Weigel, and Bruno Welsch. The CS Freiburg robotic soccer team: Reliable self-localization, multirobot sensor integration, and basic soccer skills. In M. Asada and H. Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, Lecture Notes in Artificial Intelligence, pages 93–108. Springer-Verlag, 1999.
3. Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawa, and Hitoshi Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.
4. Bernhard Nebel, Jens-Steffen Gutmann, and Wolfgang Hatzack. The CS Freiburg '99 team. In M. Veloso, E. Pagello, and H. Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*, Lecture Notes in Artificial Intelligence, pages 703–706. Springer-Verlag, 2000.
5. Thilo Weigel, Willi Auerbach, Markus Dietl, Burhard Dümmler, Jens-Steffen Gutmann, Kornel Marko, Klaus Müller, Bernhard Nebel, Boris Szerbakowski, and Maximilian Thiel. CS Freiburg: Doing the right thing in a group. In P. Stone, G. Kraetzschmar, T. Balch, and H. Kitano, editors, *RoboCup-2000: Robot Soccer World Cup IV*, Lecture Notes in Artificial Intelligence. Springer-Verlag, 2001.
6. Thilo Weigel, Alexander Kleiner, Florian Diesch, Markus Dietl, Jens-Steffen Gutmann, Bernhard Nebel, Patrick Stiegeler, and Boris Szerbakowski. CS Freiburg 2001. In P. Stone, G. Kraetzschmar, T. Balch, and H. Kitano, editors, *RoboCup-2001: Robot Soccer World Cup V*, Lecture Notes in Artificial Intelligence. Springer-Verlag, 2002.