

Brainstormers Tribots Team Description

Roland Hafner, Sascha Lange, Martin Riedmiller, and Stefan Welker

Neuroinformatics Research Group,
Institute of Computer Science and Institute of Cognitive Science,
University of Osnabrück, 49069 Osnabrück, Germany
`first_name.last_name@uos.de`

Abstract This paper describes the main reasearch activities of the robot soccer team “Brainstormers Tribots”. Improvements have been made in the field of efficient sensor fusion, self-learning robots, and stereo vision.

1 Introduction

The Middle Size League team of the *Brainstormers Tribots* was founded in 2002. Benefiting from the experiences of our simulation league team *Brainstormers*, we were able to create a new Midsized league team within two years. In 2004 we won the German Open, Europe’s major RoboCup event, for the first time in the Middle-Size League. Since then, we have successfully defended our title up to now in all three following events. At the RoboCup 2006 as well as at the RoboCup 2007 we have won both, the Middle-Size competition and the Technical Challenge. 2008 we achieved the third Place in the German Open Competitions and the third Place in the Robocup 2008 Competitions in China where we also achieved a Technical Challenge award.

The basis for our success was the robust and reliable hardware design, a well-structured software architecture and efficient algorithms for sensor fusion and behavior generation. Our main research interest is both, the development of learning robots and the development of improved sensor fusion and sensor integration techniques. Among others, several different approaches have been investigated so far:

- reinforcement learning to learn intercepting a rolling ball
- reinforcement learning to learn dribbling the ball
- learning motor commands
- development of a computationally efficient algorithm for self-localization
- development of estimation procedures for robot and ball velocity
- development of a hybrid stereo vision system integrating an omnidirectional and a perspective camera in hard real-time to recognize the ball position in full three-dimensional space
- biologically inspired real-time capable object recognition

In this paper, we will at first describe the general hard- and software design of the Tribots (section 2) and after that focus on our scientific approaches in sensor fusion and learning (section 3).



Figure 1. A robot of the Brainstormers Tribots team

2 Architecture

2.1 Hardware

Since 2004 we have been using almost the same homogenous robot hardware architecture for all the robots in the team, based on an omnidirectional mobile platform. For 2009 we started the development process for a new hardware design together with our sponsor, Harting KGaA. This development is mainly focussed on creating an open source hardware platform which is optimized for the use in the demanding RoboCup Mid-Size league games. The new hardware is again built as a triangular configuration using three omnidirectional wheels to achieve an omnidirectional behavior [14]. The advantage of this configuration is the usage of the minimum number of active wheels, that is three. Contrary to the old robots, the new platform uses powerful brushless motors and appropriate motor controllers. Also, the hardware is designed to use lithium-polymer batteries that allow for better performance and longer drive-time of the robot. The whole design is based on the principle that the operation of the robot should be easy, the hardware should be robust and the platform to be robust in order to keep manual repair efforts minimal.

To kick the ball, the new robots are equipped with a pneumatic kicking device developed by Harting KGaA. Two pneumatic cylinders allow for straight and lob shots.

On the new platform, we have either a standard laptop with 1.6 GHz CPU running Linux, alternatively a multi-core Mini-ITX based system may be used. Our software supports 2 cameras, one omni-directional, one directional.




Figure 2. New Hardware Platform for the Brainstormers Tribots team

2.2 Software Design

The software was designed to primarily meet the demands of our research concerning learning robots and to be easy-to-handle at the same time. Therefore we use a clocked control loop in a single-threaded process. This allows to model the problem of controlling the robot as a time-discrete *Markov decision process* which is the basis for reinforcement learning [20].

The software is built in a modular way using the *design-by-contract*-principle [12] as major design criterion. This helps to easily understand the structure of the program, to implement alternative realizations for subtasks, to compare different algorithmic approaches under the same basic conditions and to check the correctness of the software as well as its real-time behavior.

The behavior of the robots is generated from a modular system of individual building blocks named *behaviors*. Each behavior is targeted to reach a certain goal, e.g. dribbling of the ball to a certain target position, blocking an opponent offender or waiting for a pass. Logical conditions are used to control which behavior may become active in a certain situation. Currently, the complete behavior architecture contains fifty different building blocks which are grouped in several levels of a hierarchy realizing the strategy of attacker, defenders and the goal keeper.

To improve the general level-of-play in the RoboCup middle-size league we published the software that we used in the RoboCup German Open 2005 as well as the tools that we developed for calibration and real-time debugging under the terms of the *general public license* (GPL). Another release is scheduled for

beginning of 2009. The software package can be downloaded from our website <http://www.tribots.uos.de>.

3 Research Focus

3.1 Reinforcement Learning




Figure 3. Milestones on the Brainstormers' road towards self-learning soccer robots.

Our group's main research focus is on reinforcement learning (RL). The idea of reinforcement learning is based on the principle of interaction with a system and learning from success and failure.

We have applied RL methods to different problems in the SSL since 1998 (see figure 3 for a list of major milestones). The soccer simulator allows for simulations faster than real-time that were a necessity for applying RL-algorithms available at that time. Our main focus in the SSL learning experiments has been on benchmarking newly developed RL algorithms and on multi-agent learning [7,11,15,?,1,19].

Within our MSL team we try to bring RL to real world systems. We started by still doing the learning in a simulator and then applying the resulting, fixed strategy to the real robots [13,2]:

- In [2] we showed that it is possible to learn the behavior of approaching a ball lying on the ground by accelerating and decelerating the wheels of the robot. This task has been learned using a physical simulation of the drive unit and has been successfully transferred to the real robot afterwards. The complexity of the task arises from the complex dynamics of the system on the one hand and from the objective to learn a time-optimal policy, on the other hand. Furthermore, some constraints had to be fulfilled: (a) the robot is not allowed to touch the ball, and (b) the control policy is not allowed to exceed the maximal wheel speed of the robot. This task is therefore a combination of a control problem and a trajectory planning task in the space of wheel

speeds. Decomposing this problem in the classical way into motor controller design and trajectory finding would probably exclude optimal solutions of the entire problem.

- In a second experiment we let the robot learn to intercept a rolling ball [13]. Again, we used reinforcement learning techniques. The learning algorithm was applied to several thousands of training experiments which were gathered using a simulator. The finally learned policy was successfully applied on the physical robot.

Recently, we proposed a new neural network based RL batch-mode method that reduces the number of interactions significantly, thus allowing to learn optimal strategies directly on real world systems without any prior knowledge (or simulation) of the system dynamics. This method, named NFQ [17], allowed our middle size robot to learn a capable dribbling behavior within a few minutes of playing with the ball. We received the 2007 MSL Technical Challenge award for demonstrating this experiment in Atlanta.

Furthermore, this method has been successfully applied to other real world tasks like balancing an inverted pendulum [18], steering the robot car Junior [16] and controlling a DC motor [3]. The replacement of the classical PI control scheme for dc motors like those employed on our robots with a learned controller is still an ongoing research project. The goal is to learn an optimal, non-linear controller for a specific task and motor purely by interacting with the system for a short time. Using this approach, it is possible to overcome system anomalies like non-linear motor and power amplifier behavior and to adapt the controller to different dynamic states of the complex system the motor is part of. As a first result, using NFQ, we were able to learn an optimal control strategy for the motors of the Tribot's drive unit in less than five minutes of interaction [3].

3.2 Sensor fusion

One of the most important subtasks in RoboCup is the calculation of a reliable and consistent knowledge base of what is going on in the robot's environment, i.e. the task of building a dynamic model of the world. As sensory input we only use the images of the omnidirectional camera and the wheel encoder values, both heavily affected by noise.

This task contains several parts which have been successfully tackled in the recent years:

- **self-localization**: the robots must be able to determine their position and heading on the field. We developed an algorithm based on the principle of minimizing the distance between the white field markings found in the camera image and the expected position of the field markings due to the RoboCup rules [9]. This approach has an accuracy of below 14cm and needs only 6ms computation time per image.
- **ball motion**: to build a model of the ball velocity we use a regression approach that estimates the position and velocity of the ball from the camera

images. It assumes a linear movement of the ball but takes into account the possibility of sudden changes due to collisions or a robot kicking the ball [10]. Furthermore it is invariant to the ego-motion of the observing robot.

- **ego-motion and collisions:** Due to the physical restrictions of the robot, slippage and collisions with other objects the actual velocity of the robot differs a lot from the desired velocity as well as from the velocity measured by the wheel encoders. To overcome that problem we developed an algorithm to estimate the ego-motion of the robot using a kinematic model of the robot combined with a regression model specialized to estimate curved trajectories [8]. This velocity estimator is also used to detect collisions of the robot with other objects comparing the desired velocity with the estimated one.

To meet the demands of the larger fields used in the RoboCup MSL since 2007 we improved the sensory system of the robot with a gyroscope which yields high precision data on the turning of the robot. This enables the robot to localize more robustly even in situations in which only few white field markings can be recognized in the camera images.

3.3 Computer Vision

Advancing real-time capable computer vision methods has already been of some interest to us [5,6,4]. Lately, we have developed a hybrid stereo vision approach allowing for real-time stereoscopy at 30 fps on our middle-size robots [21]. The prototype of this system won the 2006 MSL Technical Challenge.

In addition to the already existing omnidirectional camera we mounted a second, perspective camera on the robot that is pointing to the front. While the omnidirectional camera is used as before to recognize the white field markings, the goals, the teammates, and the opponents the perspective camera is used to recognize the ball from a different point of view.

Once having found the ball in both camera images we can determine its three dimensional position by geometric reasoning finding the closest point to two skewed lines which are defined from the two camera images. To achieve information about the vertical movement of the ball we extended the ball motion estimator by a component to estimate the vertical velocity. The motion model takes into account gravity and bouncing effects.

With the start of our new Humanoid team, the Brainstormers Twobots, we want to even intensify our efforts in the computer vision domain, planning the realization of low-level image processing algorithms directly in hardware in order to increase the power efficiency and reduce the computation time significantly. The algorithms, which will be specified in VHDL and run on a custom SoC on FPGA, will work directly on the pixel stream of a CMOS chip, handling the major part of the low-level preprocessing and thus, be able to significantly reduce the work load for the CPU. Although this system is primarily targeted at our new humanoid robots, we will investigate the possibility of using this system—once finished—in the middle size league as well.




Figure 4. Laser data with superimposed camera-processed data. The data from the laser range finder was used as ground truth for the empirical evaluation of the stereo vision system’s overall precision.

4 Discussion

The research areas of the *Brainstormers Tribots* concentrate on the fields of sensor processing and fusion and learning robots. While the first domain allows to build reliable and interpretable knowledge of what is going on in the robot’s environment, the second domain allows to learn optimal behavior.

The main contributions in the domain of data processing are a new real-time capable stereo vision system for the detection of the ball, a robust and accurate method for robot self-localization, an algorithm for estimating a dynamic model of the ball movement and algorithms to determine the robot velocity and to detect collisions without the help of haptic or ultrasonic sensors.

In the domain of reinforcement learning we could present approaches to learn basic skills like driving to a certain position without touching the ball and intercepting a ball. New data-efficient algorithms like NFQ allowed for self-learning a dribble behavior on the real robot without needing a simulation or prior knowledge about the system dynamics. Several learned behaviors have been used during the competitions in 2007. Ongoing research is dealing with the problem of optimal robot control on the level of motor commands.

References

1. T. Gabel and M. Riedmiller. Learning a Partial Behavior for a competitive Soccer Agent. *Künstliche Intelligenz*, 2:18–23, 2006.
2. R. Hafner and M. Riedmiller. Reinforcement learning on an omnidirectional mobile robot. In *Proc. of 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, Las Vegas, 2003.
3. Roland Hafner and Martin Riedmiller. Neural reinforcement learning controllers for a real robot application. In *2007 IEEE International Conference on Robotics and Automation, 10-14 April*, Rome, Italy, 2007.

4. Tim Kietzmann, Sascha Lange, and Martin Riedmiller. Incremental GRLVQ: Learning relevant features for 3d object recognition. In *Neurocomputing, in press*.
5. Sascha Lange and Martin Riedmiller. Evolution of computer vision subsystems in robot navigation and image classification tasks. In *RoboCup 2004: Robot Soccer World Cup VIII*, LNCS. Springer, 2005.
6. Sascha Lange and Martin Riedmiller. Appearance based robot discrimination using eigenimages. In *RoboCup 2006: Robot Soccer World Cup X*, LNCS. Springer, 2006.
7. M. Lauer and M. Riedmiller. An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In *Proceedings of International Conference on Machine Learning, ICML '00*, pages 535–542, Stanford, CA, 2000.
8. Martin Lauer. Ego-motion estimation and collision detection for omnidirectional robots. In *RoboCup 2006: Robot Soccer World Cup X*, 2006.
9. Martin Lauer, Sascha Lange, and Martin Riedmiller. Calculating the perfect match: an efficient and accurate approach for robot self-localization. In *Robocup-2005*, 2005.
10. Martin Lauer, Sascha Lange, and Martin Riedmiller. Modeling moving objects in a dynamically changing robot application. In *KI 2005: Advances in Artificial Intelligence*, pages 291–303, 2005.
11. A. Merke and M. Riedmiller. Karlsruhe Brainstormers - A Reinforcement Learning Way to Robotic Soccer II. In A. Birk, S. Coradeschi, and S. Tadokoro, editors, *RoboCup-2001: Robot Soccer World Cup V, LNCS*, pages 322–327. Springer, 2001.
12. Bertrand Meyer. Applying "design by contract". *IEEE Computer*, 25(10):40–51, 1992.
13. Heiko Müller, Martin Lauer, R. Hafner, S. Lange, A. Merke, and M. Riedmiller. Making a robot learn to play soccer using reward and punishment. In *KI*, 2007.
14. F.G. Pin and S.M. Killough. A new family of omnidirectional and holonomic wheeled platforms for mobile robots. In *IEEE Transactions on Robotics and Automation*, pages 480–489, 1994.
15. M. Riedmiller and A. Merke. Using Machine Learning Techniques in Complex Multi-Agent Domains. In I. Stamatescu, W. Menzel, M. Richter, and U. Ratsch, editors, *Adaptivity and Learning*. Springer, 2003.
16. M. Riedmiller, M. Montemerlo, and H. Dahlkamp. Learning to drive a real car in 20 minutes. In *Proceedings of the FBIT 2007 conference, Jeju, Korea*, 2007.
17. Martin Riedmiller. Neural fitted Q iteration - first experiences with a data efficient neural reinforcement learning method. In *In Proc. of the European Conference on Machine Learning, ECML 2005*, Porto, Portugal, 2005.
18. Martin Riedmiller. Neural reinforcement learning to swing-up and balance a real pole. In *In Proc. of the Int. Conference on Systems, Man and Cybernetics 2005*, Big Island, USA, Octobre 2005.
19. Martin Riedmiller and Thomas Gabel. On experiences in a complex and competitive gaming domain: Reinforcement learning meets robocup. In *Proc. of the 2007 IEEE Symposium on Computational Intelligence and Games (CIG 2007)*, pages 17–23. IEEE Press, 2007.
20. Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
21. Arne Voigtländer, Sascha Lange, Martin Lauer, and Martin Riedmiller. Real-time 3D ball recognition using perspective and catadioptric cameras. In *Proc. of 2007 European Conference on Mobile Robots (ECMR07)*, Freiburg, 2007.