

Rhoban Football Club – Team Description Paper

Humanoid Kid-Size League, Robocup 2019 Sydney

J. Allali, L. Gondry, L. Hofer, P. Laborde-Zubieta,
O. Ly, S. N’Guyen, G. Passault, A. Pirrone, Q. Rouxel
julien.allali@labri.fr, loic.gondry@free.fr, lhofer@labri.fr,
patxi.laborde.zubieta@gmail.com, ly@labri.fr, steve.nguyen@labri.fr, g.passault@gmail.com,
antoine.pirrone@gmail.com, quentin.rouxel@labri.fr

CNRS, LaBRI, University of Bordeaux and Bordeaux INP,
33405 Talence, FRANCE

Abstract. This paper¹ presents a brief overview of the design of the Kid-size humanoid robots from the *Rhoban Football Club* Robocup team. It focuses on the description of our existing software, hardware and research projects related to playing soccer autonomously. We emphasize on the specificities of our team and on recent improvements.

1 Introduction and Last Participations

Rhoban Football Club² is an on-going robotic project whose team members are researchers and PhD students at University of Bordeaux (France), CNRS and Bordeaux INP.

The interest of the team is mainly focused on autonomous legged robots and their locomotion. Our two leading projects are a small and low cost open source quadruped robot³ and a Kid-size humanoid robot with the RoboCup competition as major ambition (Sigmaban+ platform). In this context, several prototypes have been built and tested with a special emphasis on pragmatic and operational solutions.

The very challenging problem of robots playing autonomous soccer in complex and semi-unconstrained environment has driven the team to propose new mechanical designs – spine-oriented robot have been tested, low-cost foot pressure sensors are experimented – and software methods – new custom servomotors firmware, learning algorithms applied to odometry, motion generation

¹ This document contains elements from the previous Team Description Paper, in order to make clear for reviewers which elements have been added or modified since last year, we color in blue the new or significantly rewritten content. Note that some of the new content presented in this document comes from [3], which contains more detailed information on the improvement brought to our robots between RoboCup2016 and RoboCup2017.

² The page of the team is accessible at: <http://rhoban.com/robocup2018>

³ Metabot Project: <http://metabot.cc>

and navigation problems.

Our participation to Robocup 2019, up to the qualification procedure, would be the eighth one:

2011 (Istanbul): Very first participation of the team at RoboCup competition under the name *SigmaBan Football Club*.

2013 (Eindhoven): Second participation under current name *Rhoban Football Club*. For the first time, the team was able to submit three robust humanoid robots without major hardware problem.

2014 (João Pessoa): We took a big step forward by reaching the quarter-finals and working out a robust walk engine.

2015 (Heifei): We coped pretty well with the new artificial grass and colorless field. We reached the semi-finals and took the third place of Kid-Size league.

2016 (Leipzig): Finally, we succeed to hit the first place of the Kid-Size league thanks to our versatile vision pipeline, an improved localization module through accurate odometry learning and the very beginning of high level team play strategy described in [1].

2017 (Nagoya): This year again, we managed to take the first place of the Kid-Size league, thanks to improvements to the vision system and to the general robustness of the robots, a new walk engine, better high level team play and strategy, and new tools to make development and debugging easier as described in [3].

2018 (Montreal): For the third time, we managed to take the first place of the Kid-Size league mostly thanks to improvements in the robustness of the robots and improvements in the accuracy and speed of last steps of the ball approach phase.

This short paper gives an overview of the Rhoban robots hardware and software system in its current state with an emphasis on recent upgrades with the aim to participate to Robocup 2019 in Sydney, Australia.

Commitment

The Rhoban Football Club commits to participate in RoboCup 2019 in Sydney (Australia) and to provide a referee knowledgeable of the rules of the Humanoid League.

2 Hardware Overview

The mechanical structure of the robot has not changed much since last year. We switched to only using *Sigmaban+* robots instead of a mix of our older and newer robots designs. See [2] and this year's robot specification paper for a more complete description.

The key hardware modifications we introduced recently are the following, focusing on robustness :

1. During RoboCup 2016, we broke several of the pressure sensors used in the feet of our robots. In order to improve robustness we switched to full Wheatstone bridge, 40 kg rated off the shelf load cells. This change strongly helped to reduce the burden of robotic maintenance during the competition. We also added more protection on the side of the feet to protect the pressure sensors' fragile wire connections, which were the main breaking point of the cells, thus greatly reducing the amount of broken load cells during the competition.
2. We designed an hot-swap power board allowing to switch the batteries without any transition phase, thus removing the need to reboot the robot during half-time.
3. We changed from an USB3 camera to a gigabyte ethernet camera due to interference between USB3 and Wi-Fi communication.
4. We changed the posture of the arms of the robots so that the elbows are in contact with the hips in case of a lateral fall, thus reducing the risk of breaking the shoulder motor's axle, which is at the end of the kinematic chain of the arm, and subject to a great lever effect. Little blocks of teflon are positioned at the contact points in order to limit friction. Thanks to this improvement, we didn't have to change any shoulder motors during the competition, whereas in 2017, we broke about 7 of these motors.
5. We also use a thicker piano wire on the front and the back of the robot to better absorb the frontal and dorsal falls.
6. Finally, we designed a small script to quickly check the hardware and software health of the robot, that is intended to be run on each robots before a match.

3 Vision

Our main vision system has been completely changed in 2017. The method we used before was mainly based on multiple hand-tailored OpenCV filtering, but for our purpose, this "standard" method has reached its limits in terms of complexity, robustness and maintainability. The new vision pipeline is much simpler and requires much less hand tuning. Moreover, it also appeared to be quite robust to the environments changes in luminosity and color.

Regions of interest (ROI) for the ball and the goal posts are extracted from the full image, using the robot's state (ground plane projection on the camera plane) and a kernel convolution on an Integral Image filter. These ROIs are then classified by a Convolutional Neural Network, see Fig. 1.

One key functionality to this system is the ability to quickly obtain a large quantity of labeled data. To do so, patches extracted from ROIs were uploaded on an online tagging tool⁴ accessible to the public. Tagging was made simple with

⁴ <http://rhoban.com/tag>




Fig. 1: Example of ball and goal detection

a responsive "*Google ReCaptcha*" style interface. This tool was used by several supporters from outside the team and allowed to get thousands of labeled patches in a matter of a few hours, thus freeing precious time to the team members. A consensus based approach was used to ensure the quality of the tagged data from non expert users, as well as a slight gamification of the system to motivate the users. The project is open source⁵ and the data is available once registered.

3.1 Classification

Convolutional Neural Networks (ConvNet) have become the state of the art methods in various computer vision tasks [6]. Several off-the-shelf very powerful architectures are available such as [7,8] but unfortunately none were usable in the very limited embedded computers of our robots. We thus designed our custom ConvNet using a c++ library with no external dependencies⁶. The aim of the approach was to design a minimal architecture able to classify ball and goal post patches with at least 95% accuracy.

After some hand-tuning we obtained a quite small network (cf. Fig. 2) able to classify small $16 \times 16 \times 3$ patches with good results (cf. Table 1).

	nb training/validation	Learning rate	Accuracy
Ball	7400/1500	0.013	96.8%
Posts	13500/2500	0.0425	96.92%

Table 1

⁵ <https://www.github.com/rhoban/tagger>

⁶ <https://github.com/tiny-dnn/tiny-dnn>




Fig. 2: Architecture of the reduced ConvNet used for ball classification in the RoboCup 2017

This architecture was used both for ball and goal post classification, with the small difference that the goal network has only 16 feature maps.

4 Localization

Our localization module is based on a particle filter which uses 3000 particles. It uses information from the referee, the vision module but also odometry in order to ensure a satisfying accuracy for high-level decision making.

The information provided by the referee allows to provide a reasonable idea of the position of the robot at kick-offs, drop balls or when a robot enters the field after a game stoppage. The pressure sensors allow us to obtain a satisfying odometry [9] thus allowing to reduce the exploration we use on particles and improving accuracy.

The visual features used to score the different particles are limited to the base of the goal posts and the corner of the field of play.

5 High-Level Decision Making

5.1 Finite state machines

The behaviour of the robot is designed using finite state machines. Transition between different states are based on various information such as game status, time spent in current state or information from the localization module.

Since debugging complex state machines based on all the information received by the robot is a difficult and tedious task, we can run our strategy module based on fake information. Thus, we allow to test quickly multiple situations without requiring to reproduce them in the real world. We also designed a tool named *BehaviorViewer* which can be used to monitor the state of the robot, but also to modify the current state by changing the positions of the ball, the robot or obstacles, see Fig. 3. Enabling easy and quick testing of complex features has proved to be a valuable asset during the last competition.




Fig. 3: An example of situation in the *BehaviorViewer* involving a robot playing with the ball, obstacles, trajectory and vision cone. Image from [3].

5.2 Ball Approach

Most of the game is spent with robots trying to reach a suitable position to perform a kick, more specifically, placing the ball accurately in presence of all the constraints (perception and action noise) takes a significant amount of time. Our team used learning methods based on *continuous state and action Markov decision processes* to act more efficiently. By learning a predictive motion model of the robot and computing offline optimization of the policy we manage to obtain more satisfying policies in the real world, see [5].

This year, a final simple stepping strategy has been implemented to improve the foot placement time at the end of the approach.

During the ball approach motion, the walk of the robot is usually controlled at velocity level (omnidirectional forward, lateral and turning velocity). But when the ball is close enough and lies in a predefined "reachable" area, a final accurate step is computed and applied before the robot to come to full stop and kick. This step allow the kicking foot to be directly placed with the right orientation in front of the ball.

In practice, depending on the current support foot and the target pose, either a single or two steps are actually issued. The choice of the kicking foot is determined in order to minimize the stepping distance and takes into account the stepping limits and as well as self foot collision constraints (for lateral steps).

To compute the desired final step displacement while the robot is moving, two key features are needed. First, the ball relative pose with respect to the robot must be estimated from the camera while carefully taking into account the vision processing delay. Second, the final step computation is only sent to the walk engine at the end of the current step (to be applied at the next (half) walk cycle). Because of this, the future robot pose with respect to the ball at the end of the current walk cycle must be also correctly estimated.

Note that because the robot comes to full stop after this last step, it turns out that a step distance larger than usual can be applied without jeopardizing the balance of the robot.

This procedure is actually a first work toward a proper steps planning strategy to optimize the several final steps with respect to stepping limits and possibly obstacles. A model predictive control (MPC) scheme is currently being experimented.

The current implementation assumes that the desired step kinematically computed is tracked and realized on the real robot. However, our past experiments on odometry showed that this is a very coarse approximation (position control error, foot sliding, joint backlash). The predictive displacement model studied and devised in [5] should be used here to encompass the reality gap.

5.3 Kicks strategies

Always kicking toward the center of the opposite goal is unsatisfying, especially when the ball is located nearby a corner of the field on the opponent half. It is also difficult to come up with satisfying heuristics based on geometrical approaches. Therefore, we model the choice of the kick as a Markov Decision Problem and we find the kicks minimizing the time required to score a goal using the Value Iteration algorithm [4].

This approach of the problem also allow us to include expert knowledge in the reward function in order to avoid kicking at the center of the goal where we usually find the opposite goalie.

We found that the direction of the blade of the grass had a strong effect on the distance traveled by the ball during RoboCup 2017, therefore we added a simple model of the grass to our problem and obtained two different kicking strategies, one for each half-time.

5.4 Teamplay

Having multiple robots collaborating is an important part of our current strategy, which was introduced and improved, relying on a good localisation of the robots which is a major requirement for teamwork.

Writing algorithms to coordinate all the playing robots during the game can be tedious if you use a decentralized paradigm, which was what we used to do in our previous teamwork [2]. This is why we re-centralized the decision in one robot which is the captain, and takes strategical decisions based on the information from all the robots.

The captain is elected by the simple rule of the lower id active robot, since only in-game players can communicate with the others. Each robot broadcast status packets, and the captain also broadcast an instruction packet with target positions and assignments.

Along with this change, we introduced better behaviours. First, the robots positioning at the beginning of a game and after a goal is scored are now assigned dynamically, which mean for instance that the robot which is closer to the kick-off position will be assigned to this position. Also, priorities were introduced to positions so that if robots are penalized there will always be a kick-off positionned one.

Also, the position of the ball of all robots are now clustered by the captain. This way, if there is an outlier (a robot seeing the ball somewhere all the others don't), he will be under-prioritized over the others. It is also possible that the robot selected to be the striker has no direct vision over the ball, but is still the nearest according to the other robots, which is useful if his viewpoint is obstructed.

References

1. Julien Allali, Louis Deguillaume, Rémi Fabre, Loic Gondry, Ludovic Hofer, Olivier Ly, Steve N'Guyen, Grégoire Passault, Antoine Pirrone, and Quentin Rouxel. *Rhoban Football Club: RoboCup Humanoid Kid-Size 2016 Champion Team Paper*, pages 491–502. 2017.
2. Julien Allali, Rémi Fabre, Hugo Gimbert, Loïc Gondry, Ludovic Hofer, Olivier Ly, Steve N'Guyen, Grégoire Passault, Antoine Pirrone, and Antoine Rouxel. Rhoban Football Club Team – Description Paper - Montreal 2018. *Humanoid KidSize League, Robocup 2018*, pages 1–8, 2018.
3. Julien Allali, Rémi Fabre, Loic Gondry, Ludovic Hofer, Olivier Ly, Steve N'Guyen, Grégoire Passault, Antoine Pirrone, and Quentin Rouxel. Rhoban football club: Robocup humanoid kid-size 2017 champion team paper. In *RoboCup 2017: Robot Soccer World Cup XXI*. Springer, Accepted.
4. Richard Bellman. A markovian decision process. *Journal of Mathematics and Mechanics*, pages 679–684, 1957.
5. Ludovic Hofer and Quentin Rouxel. An Operational Method Toward Efficient Walk Control Policies for Humanoid Robots. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, {ICAPS} 2017, Pittsburgh, Pennsylvania, USA, June 18-23, 2017.*, pages 489—497, 2017.
6. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
7. Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.
8. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
9. Quentin Rouxel, Gregoire Passault, Ludovic Hofer, Steve N'Guyen, and Olivier Ly. Learning the odometry on a small humanoid robot. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016.