Rhoban Football Club – Team Description Paper
Humanoid Kid-Size League, Robocup 2018 Montreal

J. Allali, R. Fabre, H. Gimbert, L. Gondry, L. Hofer,
O. Ly, S. N'Guyen, G. Passault, A. Pirrone, Q. Rouxel

julien.allali@labri.fr, remifabre1800@gmail.com, gimbert@labri.fr, loic.gondry@free.fr,
lhofer@labri.fr, 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 paper[1] 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 Club[2] 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 robot [3] 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 [9,10,8] 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 [2], 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 2018, up to the qualification procedure, would be the seventh 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 [2].

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 2018 in Montreal, Canada.

Commitment

The Rhoban Football Club commits to participate in RoboCup 2018 in Montreal (Canada) and to provide a referee knowledgeable of the rules of the Humanoid League.

2 Hardware Overview

The mechanical structure of the robot is a classic design using 20 degrees of freedom: 6 for each leg, 3 for each arm, and 2 for the head (pitch and yaw rotations). The global shape of the robot is mainly standard [4].

The main innovation of the robot is located in its feet. The feet are no longer flat but are put on the ground on top of 4 cleats at each foot corner. Only these cleats are in contact with the ground and "sink" into the artificial grass. This

[4] see the robot specification paper for a more complete description

greatly improve the stability of the robot walking on the "soft" turf.

In addition to the ground contact, each cleat is linked to a strain gauge force sensor. The whole is integrated into the foot with a piece of electronics and the sensor readings are published on the Dynamixel bus as a virtual device. This low-cost force sensor allows for computing an evaluation of the center of (vertical) pressure for each leg. This sensor is greatly useful to stabilize the static kick, the walk engine and improve the accuracy of the robot's odometry.

Additionally to the fact of developing our new platform, *Sigmaban+*[5], we introduce three significant hardware modifications, 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.
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.

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 tool[6] accessible to the public. Tagging was made simple with 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 approch 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 source[7] and the data is available once registered.

[5] See robot specification

[6] `http://rhoban.com/tagger`

[7] `https://www.github.com/rhoban/tagger`

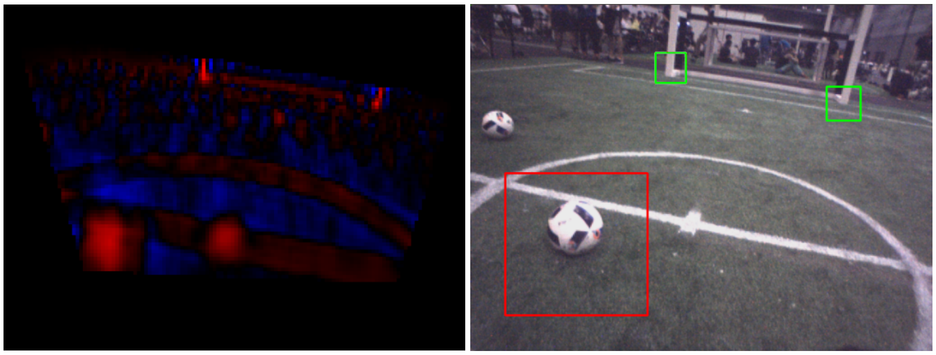

(a) Score of ROI for goal posts (b) Result of goal posts detection

Fig. 1: Example of ball and goal detection

3.1 Classification

Convolutional Neural Networks (ConvNet) have become the state of the art methods in various computer vision tasks [7]. Several off-the-shelf very powerful architectures are available such as [12,13] 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 dependencies [8]. The aim of the approach was to design a minimal architecture able to classify ball and goal post patches with at least 95% accuracy.

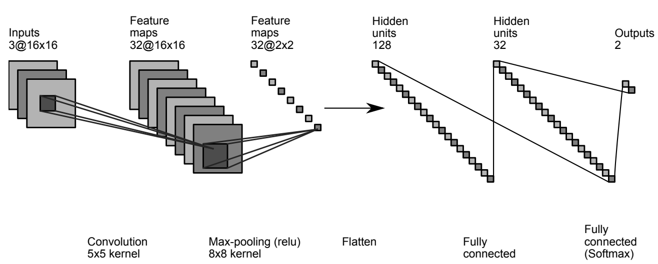

Fig. 2: Architecture of the reduced ConvNet used for ball classification in the RoboCup 2017

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).

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

[8] `https://github.com/tiny-dnn/tiny-dnn`

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

Table 1

4 Walk Engine

Last year, our walk engine has been improved to a new version, *QuinticWalk*[9], leading to a smoother and stabler motion. This new walk engine follows the same principles as the old one [14] but with many details updated.

As introduced by [3], the movement is open loop and does not use any ZMP criterion or dynamics modeling. The shapes of target trajectories are built geometrically from a set of parameters in the Cartesian space. All target joint positions are then computed through an inverse kinematics of the legs of the robot. These parameters are manually tuned on the physical humanoid robot (trials and errors) until a fast and balanced motion is achieved. The main changes are detailed in [2].

Thanks to this new walk engine, the displacement of the robot is qualitatively far stabler than before, especially on our bigger robots. For example, the lateral steps do not trigger constantly the stabilization module anymore (see [11]) which was required to prevent the robot from falling.

However, the effects of the Runge phenomenon can be observed on some parts of the trajectories. The current implementation could be improved by replacing the polynomial splines by the optimal bang-bang jerk control developed by [5].

5 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 [15] 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.

[9] Source code available at: `https://github.com/RhobanProject/Model/tree/master/QuinticWalk`

6 High-Level Decision Making

6.1 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 prooved to be a valuable asset during the last competition.

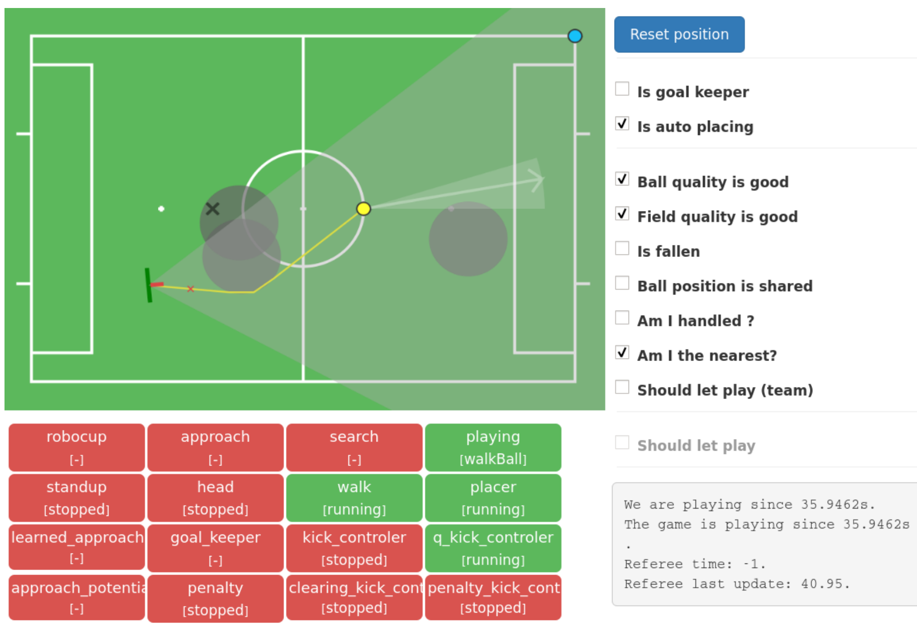

Fig. 3: An example of situation in the *BehaviorViewer* involving a robot playing with the ball, obstacles, trajectory and vision cone. Image from [2].

6.2 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 [6].

6.3 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.

6.4 Teamplay

At anytime, we consider that the robot which is closest to the ball is the kicker. It shares with other robots the position he expects the ball to be after his kick. In order to reduce the time required for the next kick, they can try to position nearby the kick target, thus making it look like a pass. On the other hand, they can also try to position between the ball and their goal in a defensive manner. We can choose how many robots play defensively or agressively by tuning up some parameters to adapt to the opponent team.

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, 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.
3. Sven Behnke. Online trajectory generation for omnidirectional biped walking. In *2006 IEEE International Conference on Robotics and Automation (ICRA 2006)*, pages 1597–1603. IEEE, 2006.

4. Richard Bellman. A markovian decision process. *Journal of Mathematics and Mechanics*, pages 679–684, 1957.
5. Xavier Broquere, Daniel Sidobre, and Ignacio Herrera-Aguilar. Soft motion trajectory planner for service manipulator robot. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, pages 2808–2813. IEEE, 2008.
6. 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.
7. 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.
8. O. Ly, , M. Lapeyre, and P.-Y. Oudeyer. Bio-inspired vertebral column, compliance and semi-passive dynamics in a lightweight humanoid robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2011)*, 2011.
9. O. Ly and P.-Y. Oudeyer. Acroban the humanoid: Compliance for stabilization and human interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010.
10. O. Ly and P.-Y. Oudeyer. Acroban the humanoid: Playful and compliant physical child-robot interaction. In *ACM SIGGRAPH'2010 Emerging Technologies, Los Angeles*, 2010.
11. Grégoire Passault, Quentin Rouxel, Ludovic Hofer, Steve N'Guyen, and Olivier Ly. Low-cost force sensors for small size humanoid robot (video contribution). In *2015 IEEE-RAS International Conference on Humanoid Robots (Humanoids 2015)*, pages 1148–1148. IEEE, 2015.
12. 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.
13. 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.
14. Quentin Rouxel, Gregoire Passault, Ludovic Hofer, Steve N'Guyen, and Olivier Ly. Rhoban hardware and software open source contributions for robocup humanoids. In *Proceedings of 10th Workshop on Humanoid Soccer Robots, IEEE-RAS Int. Conference on Humanoid Robots, Seoul, Korea*, 2015.
15. 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.