UChile Robotics Team
Team Description for RoboCup 2017

K. Lobos, G. Azócar, N. Cruz, R. Pérez, P. Miranda, F. Leiva, C. Celemín, and J. Ruiz-del-Solar

Advanced Mining Technology Center (AMTC),
Department of Electrical Engineering,
Universidad de Chile
Av. Tupper 2007, Santiago, Chile
uchilert@amtc.uchile.cl
http://uchilert.amtc.cl

Abstract. This Team Description Paper describes the organization, publications, and new developments of the UChile Robotics Team for the RoboCup Standard Platform League 2017 in Nagoya, Japan.

Keywords: RoboCup, SPL, Standard Platform League, 2017, Universidad de Chile.

1 Introduction

UChile Robotics Team (UChileRT) is a joint effort of the Advanced Mining Technology Center (AMTC) and the Department of Electrical Engineering of the Universidad de Chile in order to foster research in mobile robotics, computer vision and learning algorithms.

Our team was created in 2002 under the name of UChile1 and first participated in the Robocup, joining the Four-Legged Standard Platform League in 2003. In 2007 we changed our name to UChile Kiltros, and in 2010 we collaborated with the SPQR Italian team. After performing a complete team restructuring, due to the unsatisfactory results obtained in 2012, we were within the top twelve teams in RoboCup 2013 (The Netherlands), and we reached the fourth place in RoboCup 2014 (Brazil), RoboCup 2015 (China) and RoboCup 2016 (Germany). For the RoboCup 2017 we have developed several changes according to the new rules, mainly in vision, since the lightning on the field will not be as uniform as the previous years. In addition, several improvements have been developed in robot algorithms and strategy in order to maintain the results obtained in the last years.

This paper is organized as follows: First, in Section 2, we present past contributions from our team, followed by the developments and changes for 2017 in Section 3. In Section 4, we close by acquainting the current research lines followed by the different team members.

2 Past Relevant Work and Scientific Publications

UChileRT has been involved in RoboCup competitions since 2003 in different leagues: Four-legged 2003-2007, @Home in 2007-2015, Humanoid in 2007-2009, and Standard Platform League (SPL) in 2008-2016. UChileRT's team members have served RoboCup organization in many ways: Javier Ruiz-del-Solar was the organizing chair of the Four-Legged competition in 2007, TC member of the Four-Legged league in 2007, TC member of the @Home league in 2009, Exec Member of the @Home league between 2009 and 2015, and co-chair for the RoboCup 2010 Symposium.

Among the main scientific achievements of the group are four important RoboCup awards: RoboCup 2004 Engineering Challenge Award, RoboCup 2007 and 2008 @Home Innovation Award, and RoboCup 2015 Best Paper Award. UChile's team members have published a total of 37 papers in RoboCup Symposia (see Table 1), 27 of them directly related with robotic soccer, in addition to many papers in international journals and conferences. Finally, this year 5 papers were submitted to the RoboCup symposium.

Table 1. Presented papers in the Robocup Symposium by year

RoboCup Articles	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016
Oral	1	2	1	1	2	3	2	2	-	-	1	1	1	-
Poster	1	1	1	-	3	2	-	-	2	1	2	1	4	2

3 Developments and Changes for RoboCup 2017

3.1 Ball Perceptor improvement

In the RoboCup 2016, we presented a black and white ball perceptor based on white spots and black pentagons detection[6], which presented a high detection rate, leading us to achieve a good performance during critical matches in the SPL Competition. However, the perceptor used in 2016 presented some issues in particular situations, such as when a shadow is produced by the ball itself or when the ball is near a white spot (i.e: a line or robot part). In such situations the algorithm failed to estimate the ball's center accurately, since our color-based edges could be computed outside the ball, as well presenting false edges on the white regions. In order to address this problem, instead of using an outlier-sensitive least squares estimator to compute the center and radius, we use RANSAC due to its ability to discard edges considered as outliers. The current perceptor is now able to perform well in challenging conditions such as near the robots' feet, lines, goal posts, while keeping an operation range with distances up to 5 meters.

3.2 Multiple Hypotheses Ball Model

Since its debut in 2016, it has become evident that the new ball does not follow a linear trajectory, describing a curved path instead, due to imperfections in the manufacturing process. This imperfect ball estimator heavily hindered our game performance so we developed a new ball trajectory predictor. Our approach consist of an Interacting Multiple Model Kalman Filter (IMMKF) [9], which uses several models to predict the ball's next position, and then the degree of activation for each model is calculated according to the innovation of the Kalman filter. The final result for the ball prediction is the position and velocity given by each model, weighed by the degree of activation calculated previously. With this approach we are able to accurately predict the ball movement and achieve similar results to those with the original plastic ball.

3.3 Adaptive Vision System

Vision systems for robotic soccer have been thoroughly studied in the RoboCup due to their impact in the playing performance. As usual in any application, performance must always consider the available computational resources, so vision systems developed for specific tasks such as robotic soccer, usually remain simple and most implementation rely on color information.
In the case of the SPL, most implementations are based on or are similar to the one presented in [18], where an off-line color calibration procedure must be performed before each match, so as to classify pixels in color classes. Although such a strategy can achieve good results for playing, it is dependant on a correct calibration which requires expert knowledge, and its offline nature makes it context-specific and thus does not generalize to different environments such as different lightning conditions or carpets.

To address this issue, we propose an on-line, context-dependant estimation of the color classification (green, white, black and no-color). Since the classification step must be performed with as little operations as possible, ranges in the HSI color space are used, which are estimated online using a neural network using image statistics as inputs. The proposed system consists of multiple steps to ensure robust functionalities during matches. Since the color estimation is mostly based on green detection, we first implement a verification process to ensure estimations only on images that are likely to contain mostly the field carpet. Then a proper color range estimation is performed using the neural network models, and finally a filtering step is used to reduce instabilities in the color estimation. This work is similar to the one presented in [7], since image statistics are used to perform the color estimation, and also similar to [15], since a frame validation is used to avoid color estimation on ill-conditioned images.

Furthermore, since software alone can not achieve a completely robust vision system, we also implement a camera's parameter controller to compensate difficult and dynamic illumination conditions, by controlling the camera's exposure and gain parameters, such as the ones presented in [14] and [1], by considering image statistics and a fuzzy-controller to adapt to several points of operation, which may include spotlights (zones in the image in which the camera sensor saturates).

3.4 Efficient CNN based robot detector

Robot detection is a critical capability of robotic soccer players as it enables the players to avoid obstacles and implement complex strategies. Our previous robot perceptor was based on the one presented in [18], which achieves excellent results in controlled environments. However the static color segmentation used by this perceptor performs poorly under challenging light conditions, detecting players on shadows and spotlights.

In the past we have implemented different strategies to achieve a more robust perceptor, such as BoVW, Adaboost and convolutional neural networks, each obtaining great detections rates, but at the cost of unfeasible execution times.

To enable the use of convolution neural networks as robot perceptors in real time, this year we implement the following considerations: 1) We use the outputs of the perceptor in [18], which have high detection rate but also a high false alarm rate, as region proposals for the convolutional networks. 2) We implement and follow ad-hoc design procedures to select the hyper-parameters of the network such as the filters and number of layers, to radically reduce the inference time while preserving high accuracy. 3) We use as base fast network models such as Squeezenet[8] and XNOR-Net[16].

The implementation of the networks were performed using the Darknet [17] library, and we train the models using the dataset presented in[1]. The obtained model produces inferences on less than a milisecond on the NAO robot, with a detection rate of approximately 97 percent.

3.5 Ball Pushing Behaviors

Most of our team's reported research in the last few years has been focused on the ball dribbling, which has been modeled and trained using different methodologies, such as fuzzy controllers, reinforcement learning [11][12][13] and interactive learning [4][5].

Currently, from the ball-pushing behaviors (dribbling, kicking and passing) proposed in [11], we only use dribbling during matches, so this year we implement the remaining behaviors in our robot control system and game strategies.

To achieve this, we design these behaviors as complex Reinforcement Learning based controllers, which act over multi-dimensional action spaces (the 3 axis of the omnidirectional robotic gait). Since the complexity of models with multi-dimensional action-spaces escalates exponentially with the space dimensionality, we make use of Decentralized Reinforcement Learning, to alleviate the effects of the curse of dimensionality. We also make use of finite-support basis functions as an alternative to the Gaussian RBF state representation, to further reduce he execution times. The proposed methodology enables the use Reinforcement Learning in problems with state and action spaces, both large and multi-dimensional, using less than 0.2ms in the action selection, without sacrificing performance with respect with more standard approaches.

3.6 Whistle Detection

In former competitions, we have used the whistle detection modules developed by B-Human [19] to detect the whistle at the start of each game. This detection consists of recording the sound of a whistle in an offline procedure -the reference whistle-, and then, while operating, computing the cross-correlation between the reference and the sound recorded during the set state. Hence, if the cross-correlation outputs a value above a certain threshold, a whistle is detected. Since this algorithm is unable to generalize, a different approach is needed to detect multiple whistles.

Because of this, we have developed a whistle detection system using a statistical classifier. By means of a self-made database, a model that allows frequency-independent whistle detection has been trained. Using this approach, robots are able to generalize the whistle detection and it is no longer necessary to perform calibrations on the detector (setting up a threshold for the cross-correlation and rely on recordings).

3.7 Predicting the Ball Position for Dribbling

In the past years, an Interactive Machine Learning (IML) approach to train the dribbling behavior has been used (COACH [4][5]). The behavior obtained from using this methodology obtains great performances during both simulation and real games, but since the state-space uses only the current ball's position, it cannot account for moving balls, hence presenting a sub-optimal dribbling behavior in these situations. To address this issue (without having to modify the already trained dribbling controller or its training procedure), a new system was designed where the inputs of the controller are a modified version of those used during training, which are calculated based on the actual position of the ball. These modified inputs are obtained using information from the ball tracking algorithm, so they are calculated with the predicted position of the ball instead of the actual one (if a set of conditions are met, otherwise, the current position of the ball is used), which enables the dribbling behavior to avoid chasing a moving ball, and instead go to the ball's end position.

Early simulation results validate the proposed methodology, and show that the robot can even intercept the ball in some cases. Further validation will be performed in real matches, and we will make use of this new system during the SPL competition.

4 Current Research Lines

4.1 Reinforcement Learning

This line of work is currently being followed by two undergraduate students as part of their thesis. It is proposed to generate a methodology for implementing a decision making system, defining a state space according to specific game configurations, taking into account positions and probable team actions, and training recurrent and relevant game situations.

This work includes three main stages: *i)* the implementation or learning of tasks such as dribbling, intercepting the ball, kicking, going to strategic positions, and other similar basic behaviors; *ii)* the identification of specific game settings, and recurrent and relevant playing situations; *iii)* the reinforcement learning of high level behaviors based on a state-space transformation according to an specific game setting.

4.2 Interactive Machine Learning

This line of work is part of the doctoral thesis of one of the team members. It is proposed to develop strategies for maximizing the information obtained from users who participate in the learning process of decision making systems. Interactive Machine Learning frameworks allow learning agents to be trained by human teachers who provide different kinds of information for supporting the learning process. There are paradigms like Learning from Demonstrations (LfD) [2][3], in which the human feedback is in the actions domain, or approaches of Interactive Reinforcement Learning [10][20], in which the human feedback is in the evaluative domain.

This line of research is specially focused on methods for learning tasks of continuous actions with corrective human feedback in the actions domain. The Ball Dribbling problem has been approached with these learning methods, and currently, the training of the associated controller is based on a method proposed in this work [4][5].

4.3 Probabilistic Modeling and Data Fusion

This line of work is part of the magister thesis of one of the former team member. It is proposed to achieve an accurate tracking of robots in the field and generate a correct obstacle model. It has become apparent that a good estimation of the other robot's positions is crucial in soccer robotics. To solve this task classical approaches use a vector representation of the robot's positions and Bayesian filters to propagate them over time. However these approaches suffer from the data association problem. In order to tackle this issue, a new methodology for the robust tracking of robots using the Random Finite Sets framework was developed based on the GM-PHD filter, which does not require any explicit data association. In this approach the estimations of the robots positions and observations are represented by a Mixture of Gaussians, but instead of associating a robots hypothesis or an observation to a given Gaussian, the weight of each Gaussian maintains an estimation of the number of robots that it represents. This methodology resulted in a 25 percent reduction on the errors of the estimated robot's positions compared to the Extended Kalman filter approach.

Acknowledgments

This research is partially supported by Fondecyt project 1161500.

References

1. Albani, D., Youssef, A., Suriani, V., Nardi, D., Bloisi, D.: A deep learning approach for object recognition with nao soccer robots. 20th RoboCup International Symposium, Leipzig, Germany (2016)
2. Argall, B., Browning, B., Veloso, M.: Learning robot motion control with demonstration and advice-operators. In: n Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (September 2008)
3. Argall, B.D., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. Robotics and Autonomous Systems 57(5), 469 – 483 (2009), `http://www.sciencedirect.com/science/article/pii/S0921889008001772`
4. Celemin, C., Ruiz-del Solar, J.: Interactive learning of continuous actions from corrective advices communicated by humans. In: RoboCup Symposium 2015. Hefei, China. (2015)
5. Celemin, C., Ruiz-del Solar, J.: Teaching agents with corrective human feedback for challenging problems. In: Computational Intelligence (LA-CCI), 2016 IEEE Latin American Conference on. pp. 1–6. IEEE (2016)
6. Gabriel Azócar and Pablo Cano: UChile Team Ball Perceptor 2016. `https://github.com/uchile-robotics/nao-ball-perceptor-2016`
7. HTWK, N.T.: `https://github.com/NaoHTWK/HTWKVision`
8. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size (2017)
9. Kirubarajan, T., Bar-Shalom, Y.: Kalman filter versus imm estimator: when do we need the latter? IEEE Transactions on Aerospace and Electronic Systems (2003)
10. Knox, W.B., Stone, P.: TAMER: Training an Agent Manually via Evaluative Reinforcement. In: IEEE 7th International Conference on Development and Learning (August 2008)
11. Leottau, D.L., Celemin, C., Ruiz-del Solar, J.: Ball Dribbling for Humanoid Biped Robots: A Reinforcement Learning and Fuzzy Control Approach. In: Rob. 2014 Robot Soccer World Cup XVIII Preproceedings. Joao Pessoa, Brazil. (2014), `http://fei.edu.br/rcs/2014/RegularPapers/robocupsymposium2014\_submission\_58.pdf`
12. Leottau, D.L., Ruiz-del Solar, J.: An Accelerated Approach to Decentralized Reinforcement Learning of the Ball- Dribbling Behavior. In: Work. Knowledge, Ski. Behav. Transf. Auton. Robot. AAAI-2015 (to Appear). Austin, Texas USA (2015)
13. Leottau, D.L., Ruiz-del solar, J., MacAlpine, P., Stone, P.: A Study of Layered Learning Strategies Applied to Individual Behaviors in Robot Soccer. In: Rob. 2015 Robot Soccer World Cup XIX Preproceedings. Hafei, China (2015)
14. Neves, A.J., Trifan, A., Bernardo, C.: Self-calibration of Colormetric Parameters in Vision Systems for Autonomous Soccer Robots. In: RoboCup Symposium 2013 (2013)
15. Quian, Yonbo, L., D, D.: Adaptive Field Detection and Localization in Robot Soccer (in-press). In: RoboCup Symposium 2016. Leipzig, Germany (2016)

16. Rastegari, M., Ordonez, V., Redmon, J., Farhadi, A.: Xnor-net: Imagenet classification using binary convolutional neural networks. CoRR (2016)
17. Redmon, J.: Darknet: Open source neural networks in c. http://pjreddie.com/darknet/ (2013–2016)
18. Röfer, T., Laue, T., Kuball, J., Lübken, A., Maaß, F., Müller, J., Post, L., Richter-Klug, J., Schulz, P., Stolpmann, A., Stöwing, A., Thielke, F.: B-Human team report and code release 2016 (2016), only available online: http://www.b-human.de/downloads/publications/2016/coderelease2016.pdf
19. Röfer, T., Laue, T., Richter-Klug, J., Schünemann, M., Stiensmeier, J., Stolpmann, A., Stöwing, A., Thielke, F.: B-human team report and code release 2015 (2015), only available online: http://www.b-human.de/downloads/publications/2015/CodeRelease2015.pdf
20. Vien, N., Ertel, W., Chung, T.: Learning via human feedback in continuous state and action spaces. Applied Intelligence 39(2), 267–278 (2013), http://dx.doi.org/10.1007/s10489-012-0412-6