Team Edinferno Description Paper for RoboCup 2013 SPL

Aris Valtazanos, Efstathios Vafeias, Alejandro Bordallo Mico, Daniel Mankowitz, Nantas Nardelli, Subramanian Ramamoorthy, Sethu Vijayakumar

School of Informatics
The University of Edinburgh
10 Crichton Street
Edinburgh EH8 9AB
United Kingdom

Abstract. This paper summarizes progress made by our robotic soccer team, *Team Edinferno*, towards our participation in the 2013 RoboCup Standard Platform League competition held in Eindhoven, The Netherlands. Our team made its RoboCup debut in the 2011 world cup held in Istanbul, Turkey, where we entered both in the Standard Platform League and the 2D Simulation League. In our debut at the SPL, where we entered a mostly homegrown team, we were eliminated after the first round. We returned to the 2012 Mexico City RoboCup competition, where our team reached the quarterfinals, losing to the defending champions and eventual finalists, B-Human. This version of our code leveraged the publicly available B-Human framework to provide us a software base and modules for walking and low-level vision. We implemented and improved on all other modules including robot behaviour, optimised kicks and goalkeeper behaviour, team level coordination/communication and probabilistic localisation. For the 2013 competition, we will retain the same base software framework, allowing us to implement our recent advances in vision, self-localisation, and behaviour algorithms. As a team, our primary research interests are centered on issues of robot learning, especially for effective autonomous decision making and strategic behaviour in continually changing worlds.

1 Introduction and Team Composition

Team Edinferno is a relatively new team, consisting of graduate students and experienced researchers from the School of Informatics at the University of Edinburgh. We come from a strong research group studying robot learning, situated within a diverse community of AI and computer science researchers - the largest and best in the UK. The team leader is Dr. S. Ramamoorthy, who has extensive background in robotics and machine learning, in academia and industry. Research within our group is organized around the theme of developing autonomous decision making mechanisms in continually changing and strategically rich environments, while also leveraging our established strengths in robot control and motion synthesis.

Edinferno was one of the smallest delegations among participating teams in the 2012 RoboCup (Figure 1), where our actively involved team comprised only four members (including the team leader), in contrast to some other teams with significantly more manpower and resources. This has implied that we have had to focus our efforts on substantially advancing specific subsystems where we bring unique research expertise and competitive advantage, while leveraging a publicly available framework as a software base. Specific areas in which we bring novel solutions this year include robot behaviour (full body motion for kicks, dribbles and the goalkeeper - building on last year's work; motion strategies learnt from human demonstration [4]), localisation (using efficient features and landmarks extracted from natural background images [14]) and team level coordination (more robust tactics for use of shared team state, accommodating for system failure such as the poor wireless channels in Mexico City). Also, we are working to make our base modules more robust.

Fig. 1. 2012 Edinferno SPL team. *Left to right:* Efstathios Vafeias, Subramanian Ramamoorthy, Aris Valtazanos, Christopher Towell.

2 Basic Components

Software Architecture Our earliest attempt at robot soccer was based on a completely home-grown implementation based on Aldebaran's NaoQi. By the time we competed in RoboCup 2011, we moved to an implementation based on the rUNSWift framework[1], partly to benefit from the established rUNSWift walk, but also because our limited resources were better served by focussing our own attention on specific modules where we had the expertise and manpower to innovate, leveraging existing technology from elsewhere. So, we exploited rUNSWift's pre-existing low-level functionalities for walking, sensor management and memory access. After a season of trials and after careful consideration of our long term needs, we decided to switch to the B-Human framework [2], which gives us a faster, more energy-efficient walk. The overall structure of the B-Human framework is more intricate than with the rUNSWift one, but it also features several flexible components, such as the Extensible Agent Behaviour Specification

[1] http://www.cse.unsw.edu.au/ robocup/2010site/reports/report2010.pdf
[2] http://www.b-human.de/downloads/bhuman11_coderelease.pdf

Language (XABSL). Thus, it is possible to re-use low-level components, while focusing our own attention on the development of high-level algorithms for behaviour, etc. - our primary research focus where we expect to bring novel innovations to the league.

Probabilistic Localisation Our 2011 entry implemented a particle filter algorithm with heuristics for resampling. A drawback of this original system was with the limited number of features we could visually obtain (central circle, goalposts, penalty spots), which forced the algorithm to often make ‘blind’ updates, leading to system level unreliability. In 2012, the core algorithm was changed from a simple particle filter to a probabilistic multi-hypothesis tracker. We use a number of hypotheses to estimate a probability distribution over robot pose. To track each hypothesis we implement an Extended Kalman Filter (EKF). By maintaining a number of hypotheses, it becomes possible to exploit the advantages of EKF without having to worry about cases where the EKF breaks down, e.g., in the kidnapped robot scenario. Such an implementation requires some ‘bookkeeping’ of hypotheses, and since in each step all hypotheses should be updated it is preferable to have a few hypotheses active at each time step. This in turn means we have mechanisms for hypothesis generation, track splitting and hypothesis pruning.

The significant change in the SPL rules starting in 2012 is that both goals are yellow and disambiguating the attacking side from the defending side becomes a major issue. In 2012, we addressed this problem through team information sharing. Essentially, we used the goalkeepers model of the world as reference point to disambiguate between sides. The goalkeeper is the robot with the least distance walked in a match hence he has the lowest probability of crossing to the opponents side. At each localization step the robot used its own observation of the ball and the transmitted ball position from the goalkeeper in order to infer its position. The competence of this approach is affected by two factors, a) network stability , b) the ball being in the goalkeepers field of view. This approach performs reasonably well in cases where the ball is on our side of the field and in the field of view of our goalkeeper, but it fails to deal with mirroring situations when the ball is on the opponents side. Another problem we had to deal with is with scenarios where robots are cluttered in the center of the field, since ball positions in the center do not provide any information about the uniqueness of the observation.

To solve with the aforementioned issues we expanded our approach with two major directions, first we wanted to make the robots more autonomous so as to be able to deal better with wireless problems. The second direction was to use information fusion from all the active players in the field. To enable the robot to deal with the mirroring problem we use game state information to reset the localization hypotheses when needed. We also take into account the odometry information to minimize the generation of new hypotheses that do not match the odometry data. Information fusion is the next big step to form a more robust algorithm to solve the mirroring issue. This year we use the merged information of all the players and not just the goalkeeper, this will allow for a better field coverage something that will be crucial in the new bigger field. To cross validate our position with our teammates we use the estimated ball position from all the robots and the negative information about the absence of a ball at certain places.

Robot Behaviour and Kicking For the 2012 competition, our behaviours were revised and ported to the new XABSL scripting language. XABSL is fully decoupled from the core functionalities of the system, and as such there is no need to recompile the entire code every time a change is made. This helped us rapidly develop and test new behaviours, both prior to and during the competition. Before the 2012 RoboCup, our work was primarily focused on developing the goalkeeper behaviour, and establishing the appropriate thresholds for when this player should come out of the goal to intercept, dive to save the ball, etc. For field players, we primarily worked on disambiguation protocols, to prevent all players from going to the ball at once. However, these also turned out to be prone to problems when the underlying wireless communication channel went down, so occasionally our players collided with each other and were penalised for pushing.

The most important modification we made during the competition was the introduction of dribbling in the second round-robin. Strictly speaking, we did not so much implement a dribbling manoeuvre as we tuned and modified the walking speed on approach to and after kicking the ball in order to implicitly achieve the dribble. This approach turned out to be surprisingly successful, considering the very little amount of time we had to test it, and was instrumental in winning our closest match in that round, our 3-2 victory against Northern Bites. We have tried to build on such insights by more systematically exploring kicks and dribbles, in the hope that it will be a useful skill for those tight situations.

When testing our behaviours in Mexico City, we discovered that our original kick was not powerful enough for the field dimensions and surface of the competition. To overcome this problem, we developed an improved kick, using analysis in a MATLAB simulation to compute the optimal joint angles and velocities via numerical optimisation. Although the execution time for this optimised kick was slightly greater than for our original kick, the resulting gain in power proved to be vital in scoring crucial goals from all points on the field.

3 Some Contributions in 2013

Natural Landmark Detection The ball-based disambiguation protocol used in 2012 proved to be sensitive to communication issues and was only intended as a first pass to understand issues. In the months following the competition, a Master's project in our lab [13] focused on natural landmark-based localisation based on speeded up variants of the BRISK algorithm. The algorithms developed in this thesis were compared to the corresponding work of the rUNSWift team in natural landmark localisation, while also being evaluated in non-RoboCup scenarios, e.g. localisation in an outdoor public environment. We are trying to port this to our 2013 RoboCup code as an additional heuristic for disambiguation in localisation, which will primarily be used when communication is lost or when a robot returns from a penalised state.

Faster Field Player Behaviours RoboCup 2012 saw the first match of our team against a very strong opponent such as the B-Human team. Although *Edinferno* displayed a competent performance in those matches, the speed of the opposing team proved to be a deciding factor in the game, preventing us from executing behaviours that we knew to be proper and strategically meaningful. For 2013, we are working on speeding up some of our core game play manoeuvres in order to address this deficit. To pick one particularly illustrative example, we noticed that whenever the ball was behind our robots, they took a lot of time to turn and align behind it. To overcome this problem, we are working on a 180° dribbling motion that will allow the robot to move the ball towards the opposing goal in such a situation, without actually performing a full turn and alignment. We expect such moves to contribute to faster and more reliable field player behaviours.

Also, as discussed in the earlier section, we have realized the need for fast (i.e., quick) execution of behaviours like kicks. After lab trials, we concluded that the ability to kick at different angles could offset orientation delays such as was necessary for the robot to align and then execute a straight kick. One approach we have experimented with is to use a reinforcement learning method, developing optimised set of kicks at different angle intervals. These could be quite useful for advancing the ball position across the field when a Nao may be time or space constrained. Then, complementing our dribbling implementation, we are developing small step-sized kicks that will improve the ball control of the Nao on the field - aimed at changing the dribbling direction of motion, the aim of these kicks is to sacrifice strength in order to maximize speed and control.

Dynamic Behaviour Thresholds In 2012, we used fixed thresholds for several aspects of the goalkeeper's behaviour, such as determining when to come out of the goal to kick the ball out. In certain situations, this approach proved to be problematic, e.g. when the ball was just outside the threshold bound but no robot was around it. We are adjusting this feature by including dynamic thresholds, which will additionally account for the motion of the other players of the team.

Mixed Role Assignment As with localisation, poor wireless signal in the 2012 competition also impacted our team-level behaviours. One reason behind this problem is that our role assignment (kicker, supporting player, defender) was fully dynamic, thus being reliant on the information received from the other players. To overcome this problem, we are investigating the introduction of a few static roles, especially for defensive play, which will come into play whenever communication is perceived to be poor (e.g. when a robot does not receive messages from the Game Controller for a long period of time). Thus, we intend to avoid situations we frequently observed in 2012, e.g. all robots going to the ball at the same time, or a robot not going to the ball despite no other teammate being around it.

Efficient Search Routines Given the new increased size of the field, the accuracy of ball detection is likely to deteriorate when a player and the ball are located on opposite sides of the field. A simple search routine, such as the one we are currently using, where a robot rotates around its own axis moving its head vertically in an attempt to locate the ball, is insufficient in these more constrained settings. To overcome this problem, we are developing a collection of new heuristic search routines, combining elements of single- and multi-agent behaviours. After a robot unsuccessfully scans for the ball for a certain period of time, it attempts to move to a new location, from where a more appropriate portion of the field may be scanned. These locations will be represented as attractor points, defined in terms of the geometry of the field. Furthermore, we are investigating coordination mechanisms for this search procedure, by which robots can decide how to best select scanning locations in order to maximise coverage of the field.

4 Future Developments and Connections to Research

Our team is composed primarily of researchers interested in intelligent autonomous robotics. So, in addition to the thrill of adventurous competition, we participate in RoboCup to advance our scientific agenda. Here, we outline a few key areas where this exchange is working effectively, between our RoboCup team and broader research.

- Strategic Interaction: We are exploring ways to make strategic decisions in response to opponent strategies that may not be known ahead of time. With this in mind, we have been investigating ways to infer the behaviour of other players in terms of pre-computed models [7], estimated finite-state models [5], distributions over template plans [6],[4], learning from human demonstration [15] etc. The goal of all of these experiments is to achieve a degree of flexibility in interactions within open environments providing limited prior knowledge.
- Multi-agent Learning in *Ad Hoc* Teams: Extending the above theme, we are investigating algorithms for multi-agent learning in *ad hoc* team settings, i.e., without prior coordination [1], [2] or models of opponent strategies [3]. Our approach is based on the use of abstractions to make game-theoretic learning algorithms tractable in problem settings of interest to robotics. While a lot of our progress so far has been in machine learning algorithms and not all of it has been translated yet to robotic implementation, we will be expressing our interest in this direction by entering at least a nominal player in the Drop-In challenge.
- Full-body Humanoid Robot Behaviours: A long standing strength within our research group is in the area of machine learning for humanoid locomotion [11] and full-body humanoid behaviours [9][3] [10] [12].
  As already mentioned above, we have started to bring in some of these ideas into our RoboCup code base. Our kick engine at RoboCup 2012 was based on simulations where we explored the space via numerical optimization and then ported promising solutions to the physical robot. This was key to our effectiveness, especially when we played against teams that were slightly faster in terms of speed of execution - our ability to control long range kicks allowed us to compensate and gain valuable space. Similarly, our goal keeper behaviour is informed by our research on full body motion synthesis.

5 Conclusions

Team Edinferno is a relatively new SPL team, the only entry from the United Kingdom. The underlying research work builds on strong background in robot learning and aims to advance the state of the art of autonomous decision making in continually changing worlds. Although the team is still in early stages, it has already achieved a performance level comparable to the league's most established teams. Building on our acquired experience, we look forward to match and, if possible, better our 2012 quarter-final result.

References

1. S. Albrecht, S. Ramamoorthy, Comparative evaluation of MAL algorithms in a diverse set of ad hoc team problems. In Proc. International Conference on Autonomous Agents and Multiagent Systems, 2012.

[3] This work was considered as a finalist for the *RoboCup Best Paper Award* at IROS 2010.

2. S. Albrecht, S. Ramamoorthy, A game-theoretic model and best-response learning method for ad hoc coordination in multiagent systems (Extended Abstract), International Conference on Autonomous Agents and Multiagent Systems, 2013.
3. M.M.H. Mahmud, S. Ramamoorthy, Learning in non-stationary MDPs as transfer learning (Extended Abstract), International Conference on Autonomous Agents and Multiagent Systems, 2013.
4. A. Valtazanos, S. Ramamoorthy, Bayesian Interaction Shaping: Learning to influence strategic interactions in mixed robotic domains, In Proc. International Conference on Autonomous Agents and Multiagent Systems, 2013.
5. M. Leonetti, L. Iocchi, S. Ramamoorthy, Induction and learning of finite-state controllers from simulation (Extended Abstract). In Proc. International Conference on Autonomous Agents and Multi-agent Systems, 2012.
6. A. Valtazanos, S. Ramamoorthy, Intent inference and strategic escape in multi-robot games with physical limitations and uncertainty, In Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems, 2011.
7. A. Valtazanos, S. Ramamoorthy, Online motion planning for multi-robot interaction using composable reachable sets. In Proc. RoboCup International Symposium, 2011.
8. A. Valtazanos, S. Ramamoorthy, NaOISIS: A 3-D behavioural simulator for the NAO humanoid robot. In Proc. RoboCup International Symposium, 2011.
9. I. Havoutis, S.Ramamoorthy, Constrained geodesic trajectory generation on learnt skill manifolds. In *Proc. Int. Conf. Intelligent Robots and Systems*, 2010.
10. I. Havoutis, S. Ramamoorthy, Motion planning and reactive control on learnt skill manifolds, *International Journal of Robotics Research*, 2013. (To Appear)
11. S. Ramamoorthy, B.J. Kuipers, Trajectory generation for dynamic bipedal walking through qualitative model based manifold learning. In *Proc. Int. Conf. Robotics and Automation*, pp. 359-366, 2008.
12. I. Havoutis, S. Ramamoorthy, Motion synthesis through randomized exploration on submanifolds in configuration space. In J. Baltes et al. (Eds.): RoboCup 2009, Lecture Notes in Artificial Intelligence, Volume 5949, pp. 92-103. Springer Verlag, Heidelberg (2010).
13. D. Mankowitz, BRISK-based Visual Landmark Localisation using Nao Humanoid Robots, *MSc Thesis, University of Edinburgh*, 2012.
14. D.J. Mankowitz, S. Ramamoorthy, BRISK-based visual feature extraction for resource constrained robots, In Proc. RoboCup International Symposium, 2013.
15. A. Valtazanos, S. Ramamoorthy, Evaluating the effects of limited perception on interactive decisions in mixed robotic environments, In Proc. ACM/IEEE International Conference on Human-Robot Interaction, 2013.