Robocup Standard Platform League - rUNSWift 2012 Innovations

Sean Harris, Peter Anderson, Belinda Teh,
Youssef Hunter, Roger Liu, Bernhard Hengst
Ritwik Roy, Sam Li, Carl Chatfield
University of New South Wales, Australia
{sharris,pjan520,tehbelinda,yhunter,rogerl,bernhardh
rdro117,samli,cchatfield}@cse.unsw.edu.au

Abstract

Robotic competitions encourage a developmental style of research and development where large scale robotic systems are incrementally constructed as a whole. This differs from the typical research approach of solving a specific problem in isolation, but is a crucial part of reaching the long-term goals of complex AI systems. This paper outlines the innovation and development of the autonomous UNSW multi-robotic system (rUNSWift) that was entered in the Standard Platform Soccer League at the International RoboCup competition in 2012. The challenge is to deliver real-time functionality within the limited resources of an on-board processor. Novel developments in 2012 include: SLAM using one-dimensional SURF features with visual-odometry as a by-product; extending foveated imaging to field-line detection; a unified field-feature sensor model; a dual-mode Kalman filter to help disambiguate the symmetric field; robot-detection data-fusing visual and sonar observations; multi-robot tracking of the ball; and omni-directional kicking. The rUNSWift system was ranked in the top three world-wide.

1 Introduction

The focus in AI on solving specific problems in isolation means that often system integration and development go unnoticed. As noted by Konforti[Konforti, 2006], many intelligent systems capable of solving complex problems are often the sum of many simpler systems. Each of these simpler systems may be unintelligent in isolation, but when properly combined they are capable of intelligent behaviours. Cohen[Cohen, 2005] also suggests that the traditional "divide and conquer" approach to research isn't the best approach to the primary long-term goals of AI. Instead he urges that the most challenging problems are best solved using a developmental approach, where an integrated whole system evolves over time. With this in mind, this paper presents the 2012 evolution of the autonomous UNSW multi-robot system (rUNSWift) that competes in the International RoboCup Standard Platform League (SPL) soccer competition (Figure 1).

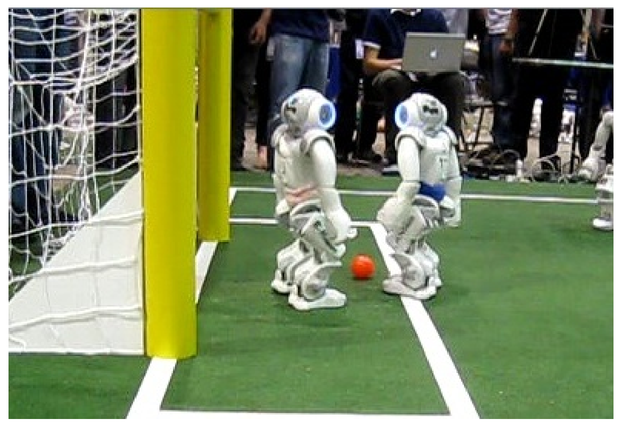

Figure 1: rUNSWift (Blue) in action at the 2012 Robocup SPL Competition.

The RoboCup Standard Platform League uses the Aldebaran-Robotics humanoid robot called Nao (Figure 2[Aldebaraan-Robotics, 2012b]). Each game runs for two 10 minute halves where two robotic teams, comprising four Naos each, compete on a 6m $\times$ 4m soccer field. In 2012, for the first time, this field is symmetric about the half-way line as both goal posts are now coloured yellow. One of the new challenges is to maintain localisation of the team during the tussle of the match to continue kicking towards the right goal. In addition, 2012 hardware changes include a processor upgraded from a 500MHz GEODE to a 1.6GHz Intel Atom, and improved resolution for both cameras that can now be accessed simultaneously, increasing the effective vertical field of view from 34 degrees to 86 degrees.

The major components of the rUNSWift robotic architecture include perception, localisation, motion, and

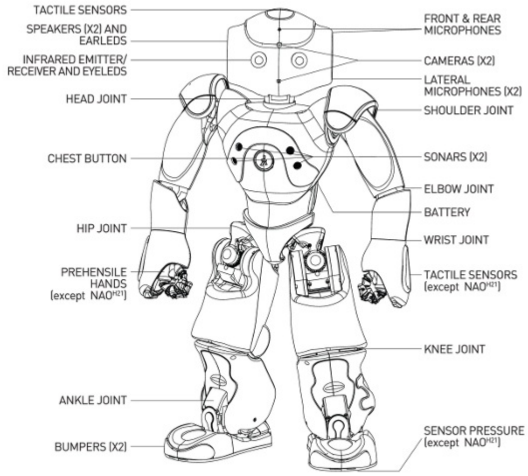

Figure 2: An Aldebaran Nao used in the Robocup SPL Competition.

behaviour [Ashar *et al.*, 2010]. The past 12 months saw considerable innovations in all components, which we describe in the rest of this paper. They are:

- Exploiting the camera resolution and increase in field of view, particularly for field-line detection using foveate imaging.
- A unified field-feature sensor model that adapts the iterative closest point algorithm to work with disparate higher level visual features. Features include the team ball to help disambiguate the field-ends.
- A "Natural Landmark" simultaneous localisation and mapping system that achieves real-time recognition by adapting SURF to one dimension and uses the special properties of horizon pixels.
- Visual odometry that uses the above speeded up robust features
- Robot Detection using a combined vision and sonar system able to detect multiple robots at varying distances and headings
- Team Ball Tracking by fusing data from multiple robots
- Multi-Modal Kalman filter to help solve the "kidnapped robot" problem and in particular to cope with the field symmetry.
- Omni-directional kicking for increased reactiveness during play.

We conclude with results showing performance improvement of individual and team behaviour that integrates these developments in the rUNSWift architecture. We also provide results showing the performance of the rUNSWift team in international competition.

2 Exploiting the Camera Resolution and FOV

The hardware improvements to the robots in 2012 offered a huge amount of potential to improve all aspects of the rUNSWift vision system. Access to both cameras at the same time allowed simultaneous tracking of the ball and field features. This in turn removed the need for any pause in game play to actively localise, since the robot is constantly receiving confirmation of its position. Figure 3 shows the extra features on offer when lining up the ball with two cameras instead of one.

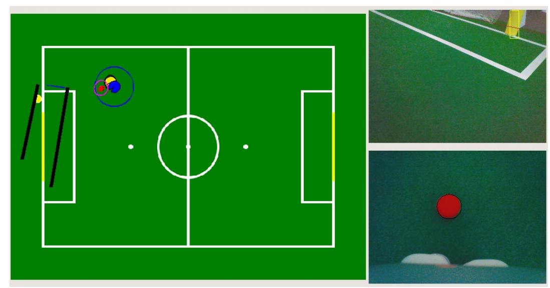

Figure 3: Left: the robot's position on the field and observed features. Right: What the robot sees whilst shooting a goal.

In addition to this, the increased processing power allowed the resolution of the top camera image to be doubled. As a result of this, the robot is able to detect field lines up to 3m away, instead of only 1.5m away in 2011. It is also able to reliably track the ball 6m (length of the field) away instead of 4.5m in 2011. Figure 4 shows the goal keeper now being able to detect the centre circle from its own goal box, allowing it to stay localised without looking away from the ball. It also shows that the goal keeper can track the ball from the other side of the field, which helps the team maintain a strong belief about where the ball is at any point in time.

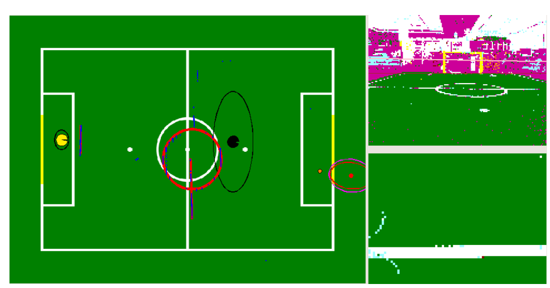

Figure 4: Left: the field showing the goal keeper's position and observations. Right: the camera images the robot can see.

3 Unified Field-feature Sensor Model

The rUNSWift vision system returns a list of field features observed by the robot each cycle. These features include goal posts, single lines, parallel lines, corners, T-junctions, circles and field edges. All of these observed features can assist with determining the robot's pose (location on the field and heading), however the nature of a soccer field means that these features aren't unique. This makes it often impossible to determine a robot's exact pose from a single feature. In saying that though, often multiple features are observed at once, providing an opportunity to combine these observations to narrow down the robot's possible pose.

In the past rUNSWift has performed multiple updates of a filter per cycle, one update per feature observed. This not only breaks the assumption of one update per cycle, but also doesn't utilise the extra information available from the specific combination of features observed. In 2012, a new approach was taken where all the observed features are combined using a modified ICP [Chen and Medioni, 1991] algorithm into one pose update, thus maximising the information utilisation of each frame.

The steps used to calculate the pose of the robot are:

1. **Mapping** observed field-features from robot relative coordinates to field coordinates, using the prior robot pose estimate x_{k-1}, y_{k-1} and θ_{k-1}. In the first iteration this robot pose estimate is provided by the localisation filter, in subsequent iterations the computed pose is used.

2. **Hierarchically Matching** observed field-features with the nearest equivalent features on the SPL field map, starting with the most distinctive feature in the early iterations, and gradually including less distinctive features in subsequent iterations.

3. **Representing** matched field-features using 2D point pairs.

4. **Weighting** the corresponding pairs appropriately, given the confidence level of the observation.

5. **Assigning** a squared Euclidean distance error metric based on the distance between point pairs.

6. **Minimising** the error metric by finding a new robot post estimate x_k, y_k and θ_k.

7. **Iterating** the above steps to further improve the robot pose and to allow any incorrectly associated features to be matched correctly.

It should be obvious that this algorithm has a lot of similarities to the standard ICP approach except that it has been tweaked for this specific application. Figure 5 shows some examples of the algorithm's ability to converge on the correct location given up to 45^o error in the initial estimate. In each diagram, the two pictures on the right are the colour classified images from the camera, essentially what the robot can see in that frame. The field view on the left shows the pose estimate (yellow) and the identified features nearby. The orange area encompasses the space in which the robot's initial pose can start and still return the correct pose estimate for that frame.

Figure 5: Some examples of the sensor model converging to the correct estimate from various locations around the field.

4 Visual odometry

A new module developed in 2012 was a visual odometry system. This vision module analyses the pixels along the horizon (pixels level with the robot's eyes) and monitors their change between frames. The feature identification is done using a novel 1D SURF developed by Peter Anderson[Anderson, 2012] that was highly optimised to run on the robot. The features matching is done using nearest neighbour matching in combination with RANSAC. The relative effectiveness at feature matching is shown in the ROC curve in Figure 6.

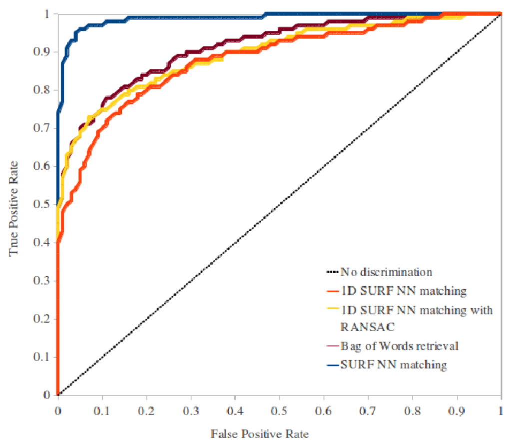

Figure 6: Comparison between SURF and Anderson's 1D SURF

Using this information, a rotational visual odometry module was created that allows the robot to track changes in its heading. The data from this module is utilised in two different ways in the rUNSWift system, one for tuning motion odometry and the other for collision detection.

The motion odometry of the robot is fairly accurate in its ability to calculate forward movement, however often struggles with the angle due to slipping of the feet when turning. By combining the visual and motion odometry readings, the system can use the best aspects of both to come up with a more accurate estimate of the robot's motion. Figure 7 shows the odometry tracks for a robot walking in a 2m box shape (clockwise and anticlockwise) around an SPL field. It highlights the difference between the original motion odometry and the improvement achieved by integrating visual features as well.

The other application of this module is collision detection. In robot soccer there are regular collisions between robots as they tussle for the ball and for position, however detecting and responding to these collisions is difficult. One interesting aspect of these collisions is that the robot is usually rotated around as a result of the contact, as opposed to being pushed backwards or forwards. This means that a significant difference between the attempted walk and the actual observed turn of the robot can be utilised to detect collisions. Figure 8 shows the difference between the motion and visual odometry when a robot spinning on the spot collides with an obstacle and is stopped from rotating.

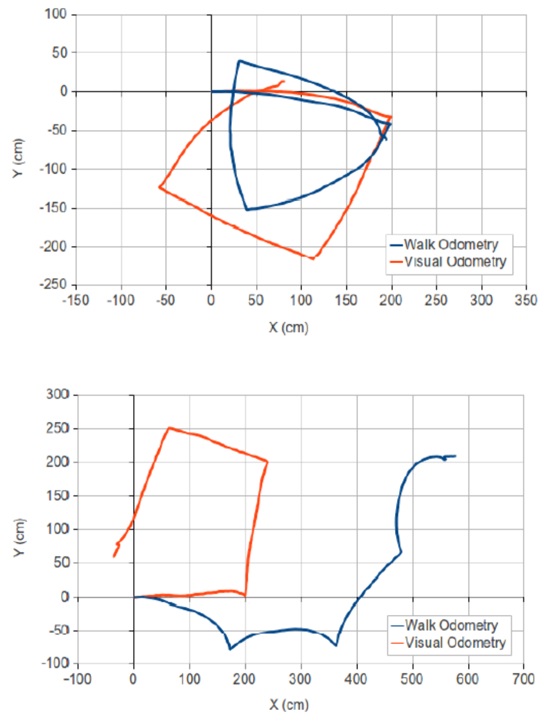

Figure 7: The difference between the raw walk odometry and the combined visual and motion odometry

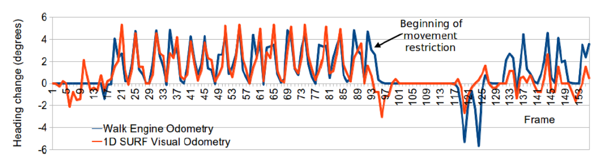

Figure 8: A graph showing how the visual odometry system agregates to zero turn once the robot becomes stuck, whils the walk odometry fails to cope

As a result of this collision detection, rUNSWift was able to limp the appropriate shoulder joint on a robot when it was caught on an opponent. This allowed the arm to slide out of the way and the robot to continue past without spinning around and falling over, a far more desirable result for the overall system.

5 Natural Landmark Localisation

As a result of the identical goal colours introduced in 2012, the SPL field no longer has any artificial landmarks to distinguish which end of the field belongs to which team. This forced the teams to come up with measures to differentiate the two field ends to allow a kidnapped robot to relocalise. The primary approach taken by rUNSWift was to use the 1D SURF features that are also utilised by visual odometry to map the horizon at either end of the field.

The process for this mapping was simple, at the start of each half the robot knows which half it starts in, so has a reliable pose estimate. From this initial position, the robot can walk to its kick-off position and turn 360^o on the spot. During this time it stores a map of the pixels behind each goal post for future matching. Once the game then starts the robot can match the mapped pixels with the currently observed pixels to differentiate the goal posts.

Storing and matching a large number of visual features can be a slow process though, so much work was done to optimise the process. The set of 1D SURF descriptors is stored in a database and the matching is done using a bag of words methodology. The features are also geometrically verified to combat the lack of ordering in the bag of words methodology. The matching is also filtered over time using a voting system and sliding window. This means that one badly matched frame won't cause the robot to turn around and shoot at the wrong goal. Figure 9 shows the process of classifying goal posts from an image.

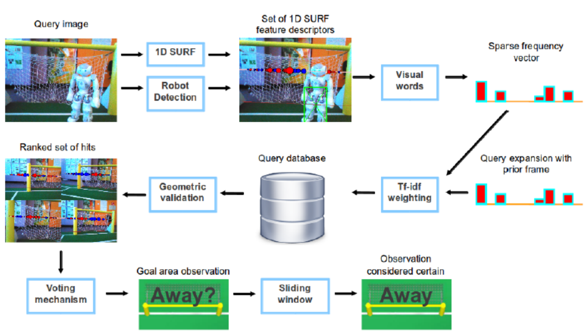

Figure 9: Flow chart showing the process to chain from image to goal post matching

To demonstrate the effectiveness of this work, the robot was kidnapped to a variety of locations around the field and then told to return to the location labelled position 1 in Figure 10. From the 10 locations, the robot was able to return to the correct location every time with an average error of 44mm. An interesting note is that in trial 5 the robot initially walked to the symmetrically opposite location to its intended destination, but as it arrived realised the mistake and corrected itself.

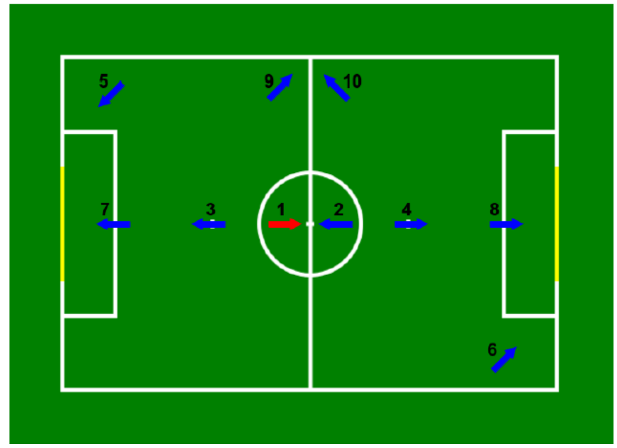

Figure 10: The experiment setup of 10 locations around the field, with the goal to return to location 1

Another experiment was set up involving a robot standing at the centre of the field and mapping the area behind the goals. Once this was done, the robot was then placed at a variety of locations around the field and the matching scores for each goal area recorded. As shown in Figure 11, the matching is good at identifying the correct goal from a similar place to where it was taken, but breaks down with large changes in angle and distance. It was particularly encouraging though to find the lack of false positive matching for the wrong goals, which contributes to the system's reliability.

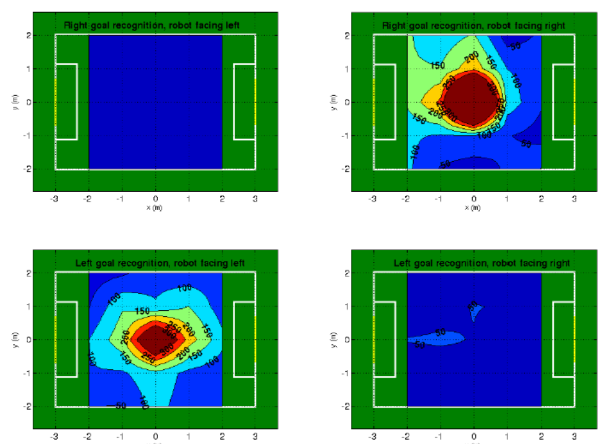

Figure 11: Heatmaps showing the matching scores for different positions around the field to each goal posts

This system was also entered to the SPL Open Challenge at the 2012 competition and placed 2nd.

6 Robot Detection Using Vision and Sonar

Robot detection for rUNSWift involves fusing a combination of visual and sonar data to detect robots. The vision side of the robot detection is based around a module from the 2011 team [Kurniawan, 2011] that utilises dips in the field edges to detect candidate robot regions (Figure 12). The original algorithm used a decision tree to then discard bad regions, but this year it was simplified to simpler checks for colours inside the region. Good candidate regions are then confirmed using sonar data to remove any false positives from bad images.

Figure 12: Visual robot detection using dips in the field edge

To enable more accurate robot detection in 2012, a new sonar filtering system was developed. The new filter has a higher directional resolution than in the past due to a unique utilisation of the sensor setup. The sonar system on the Nao is split into a left transmitter and receiver as well as a right transmitter and receiver as shown in Figure 13[Aldebaraan-Robotics, 2012a]. Obstacles can be detected on the left by turning on the left transmitter and receiver and leaving the right ones off, whilst obstacles on the right can be detected using the inverse. A new feature for 2012 is detecting obstacles straight ahead by turning on one side's transmitter and the other side's receiver. This allowed for 3 rough directions of obstacles to be identified from sonar data.

The robot detection module combines the distance and rough directional data from the sonar with the same information from the visual robot detection. The visual detector struggles to give accurate distances but is much better at determining the heading of a robot than the sonar, so the combination of the two is used to provide the most accurate data possible. The data is also filter over time to ensure that one false reading isn't enough to break the system.

To demonstrate the effectiveness of the robot detection system a penalty shootout experiment was set up. This

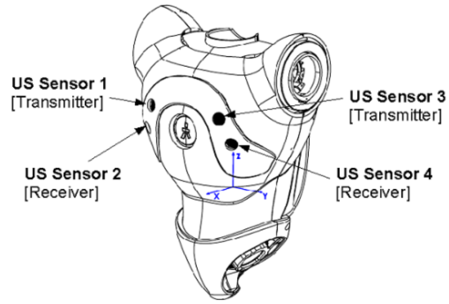

Figure 13: The sonar setup on the Aldebaran Nao

involves the ball being placed on the penalty spot, the striker on the halfway line and the goal keeper on his goal line. The goal keeper was put in three different positions and set up as shown in Figure 14. As striker walks up to the ball and shoots for a goal, the observed position of the goal keeper robot was logged. The results are shown in Figure 15. As you can see, the initial estimate of the goal keeper's pose is fairly inaccurate, but as the striker gets closer, the position becomes more and more refined. By the time the striker reached the ball, the position error of the goal keeper was between 28mm and 60mm.

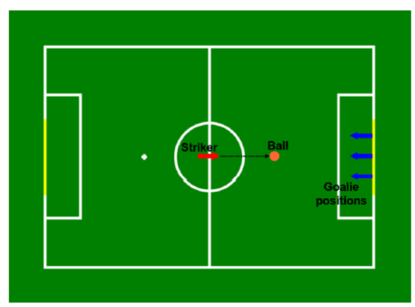

Figure 14: Robot detection experiment setup

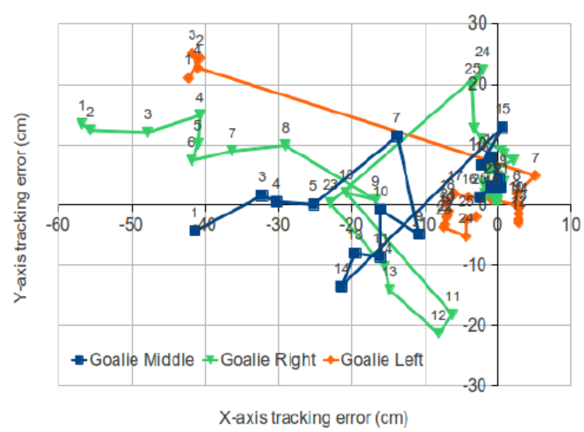

Figure 15: The plot of the goal keeper's position according to the striker as it lines up its penalty shot

7 Team Ball Tracking

Individual ball tracking is obviously an important part of playing soccer; it is difficult to make any sensible decision without knowing its location. Each rUNSWift robot tracks the ball they see using an Extended Kalman Filter with an elliptical covariance (Figure 16). In addition to each individual robot tracking their own version of the ball, it is useful for the team as a whole to know the location of the ball. Some examples of this include a striker being able to continue walking towards an occluded ball because the rest of the team can see it, or a robot realising it is mislocatised because its observed ball doesn't match the rest of the team's belief about the ball position.

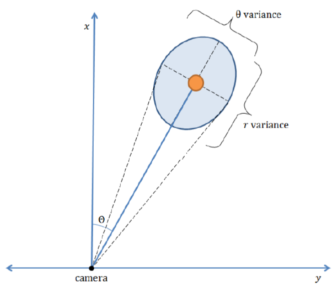

Figure 16: Individual Ball Covariance

The team ball tracking uses a linear Kalman Filter[Kalman, 1960] to track the location of the team ball. The update for the filter is the best combination of the observed balls, which might not include all the team's observations. To determine which observations agree with each other, the Malahanobis distance[Mahalanobis, 1936] is utilised in addition to a maximal absolute position difference. The observations are grouped into agreeing subsets and the best one of these subsets is chosen as the final subset. The filter update is a simple weighted average of this set of observations.

One of the useful things about the elliptical covariance is that it models the uncertainy of the distance and heading separately. This is useful since a robot is able to judge the heading of the ball quite accurately, whilst its distance is far less accurate. When the team ball combines these covariances, the ellipses often combine in a useful fashion to reduce a large portion of the uncertainty. This becomes particularly obvious when the covariance ellipses are approximately 90^o offset from each other. Figure 17 shows a good example of this.

This system is robust to noisy ball or pose estimates by individual robots as their observations are left out of the update. There is also a feedback mechanism in place where each robot is notified if its ball observation was the same as the rest of the team, giving the robot a clue that it may be mislocalised. In fact, a robot is able

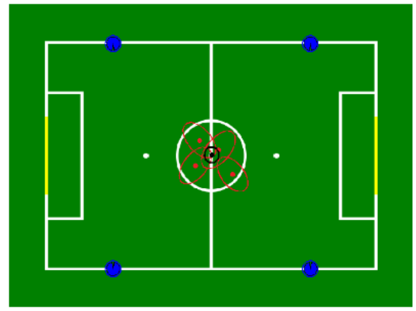

Figure 17: Team ball covariances combining to reduce uncertainy and improve the ball position

to recover from a flip to the symmetrically corresponding mode on the other side of the field if the team ball position indicates this has happened.

To show the usefulness of the team ball, an experiment was run where 4 robots stood on the sidelines of the field and observed the ball at 9 different locations around the field. This setup and the ball locations are show in Figure 18. The accuracy of the team ball improves significantly as more robots contribute to its position, as shown by Figure 19. The position of the team ball becomes particularly accurate when robots view the ball from different angles around the field simultaneously.

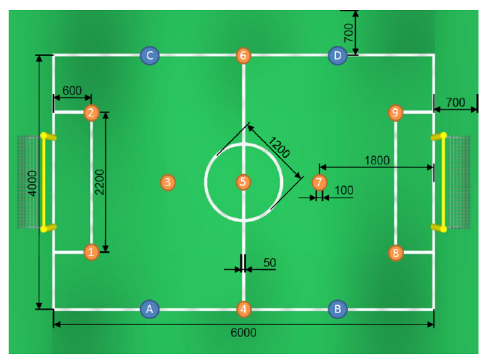

Figure 18: Experimental setup for testing team ball

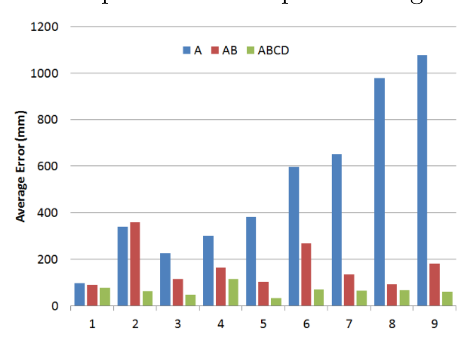

Figure 19: A graph showing how the average error is significantly reduced with more robots

8 Dual Modal Kalman Filter

The rUNSWift localisation system involves the use of a dual modal Kalman filtering approach. The actual implementation contains two unimodal Kalman filters, the "main" mode and the "alternate" mode [Hunter, 2012]. As its name suggests, the "main" mode contains the robot's current best hypothesis for where it is, whilst the "alternate" mode is the best alternative location that the robot might be in.

In order to decide how to apply updates to these two modes, the Malahanobis distance is again used. If an observation of the robot's current pose is received and it is considered similar to the main mode, then a standard update is performed. If the observation is considered reliable enough, we will also reset the alternate mode at this point as well.

However if the observation is too far away from the main mode, we consider updating the alternate mode. If the current observation matches the alternate mode well, we update it. However if it doesn't match the alternate mode either, then we need to consider which is more likely, the observation or the alternate mode. If the observation is deemed better, we can replace the alternate estimate with the new observation, else we can simply weaken the current alternate and wait for more updates.

These two filters are constantly running in parallel and the two can be switched at any point. If the alternate mode receives a significant number of updates in a row and the main mode does not, the alternate mode switches to become the main mode. One of the interesting cases here is that most updates are correct for two symmetrically opposite points on the field. The way we disambiguate the robot's heading is using the team ball and natural landmarks information described in previous sections. In this case, the alternate mode becomes the symmetrically opposite position on the field and after enough updates, the robot will "flip" to the other side of the field.

To measure the success of the system, some statistics were gathered about the robots during the 2012 competition. The first table (Table 1) shows the number of correct "Ready Skills" (Figure 20) that occured. This skill runs at the start of each half and after each goal and is where the robots positions themselves in their own half ready for a kick off. Note that the table includes an Error column for other cases where the robot didn't reach its destination due to external events including things like being penalised, falling over, running out of time, etc.

Another measure of the robot's abiliity to localise at competition was how accurately it shot during games. Table 2 shows the accuracy of the kicks made by rUNSWift during the 2012 competition. A correct kick

Figure 20: rUNSWift with a successful Ready Skill

	Localised		**Mislocalised**	
Game	Correct	Error	Incorrect	Flipped
RoboCanes	19	0	1	2
Dutch Nao Team	39	0	0	1
B-Human	19	4	0	0
TJArk	17	3	2	1
Austrian Kangaroos	19	5	0	2
Austin Villa	51	2	0	1
HTWK	49	2	0	2
Total	213	16	3	9

Table 1: Ready skill performance results from official matches

was one aimed directly into the opponent's goal while a close kick was one that narrowly missed the goal or hit the post, suggesting a small error in the pose estimate. A missed kick was one with a significant error where the ball was kicked over the sideline and finally a flipped shot was one facing our own goal.

Match (versus)	Correct	Close	Miss	Flipped
RoboCanes	17	1	0	0
Dutch Nao Team	19	0	1	0
B-Human	15	0	1	0
TJArk	15	2	1	3
Austrian Kangaroos	18	0	1	1
Austin Villa	26	1	1	0
HTWK	24	4	1	1
Total	134	8	6	5

Table 2: Kick accuracy results from official matches

9 Omni-directional Kicking

Most teams in the SPL shoot for the goals by walking behind the ball and kicking straight. rUNSWift developed a unique approach in 2012 where the robot simply walks up near the ball and turns to face the goal as part of the actual kick motion [Teh, 2012]. The method for doing this is to add an extra step phase at the start of the kick where the foot is placed at an angle parallel with the line from the ball to the goals. This is similar to how real soccer players kick the ball and allows flexibility in the shoot direction without comprimising power. Figure 21 shows the kick in action, starting with the angled step and finishing off with a big kick.

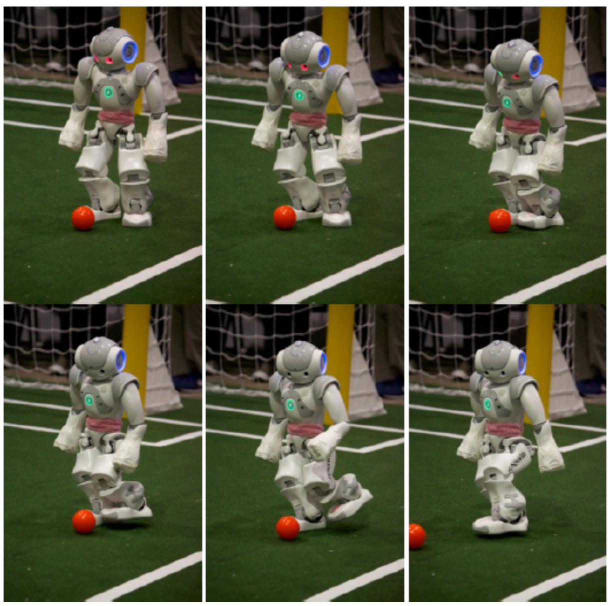

Figure 21: The rUNSWift kick in action

In addition to improving the robustness of the kicking, work was also focused on speeding up the kick action as much as possible. A faster kick gives the opponent less time to steal the ball away, so is an important aspect of soccer playing. Figure 22 shows the improvement to the speed of the kick action over stages of development this year. An interesting point to notice is that as well as being faster on average, the final result is also a lot more consistent and reliable than past kicks.

As a result of the increased flexibility in kicking directions, we were able to significantly reduce the time taken for the robot to line up and shoot for the goal. Table 3 shows the time taken to score a goal from three given starting positions around the field. This is a standard test that we performed reguarly throughout development for the past 3 years and the table shows the best time for each position each year. It's worth noting that the kicking isn't the only influence on these times, however

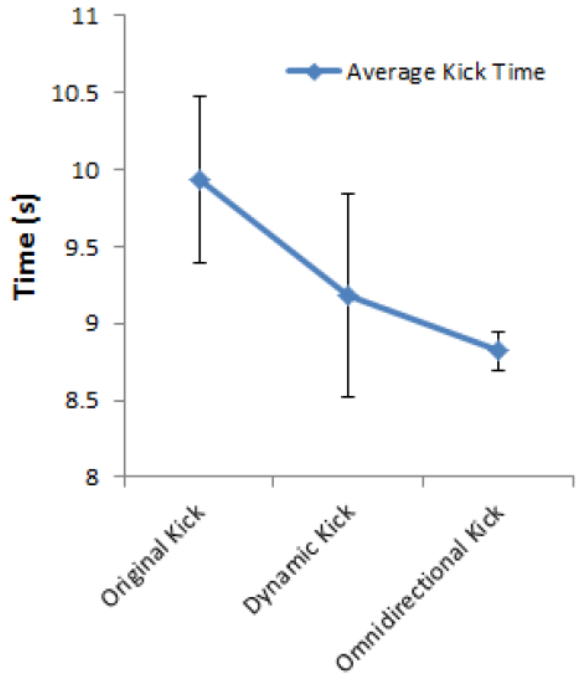

Figure 22: The improvements to the kicking time through the development cycle

was a major step in reducing the amount of time taken to shoot.

Start Position	2010	2011	2012
a) -1200, 0, 0	00:13	00:16	00:11.04
b) -1200, 2000, $-\pi/2$	00:34	00:31	00:19.00
c) -1200, 0, $-\pi$	00:40	00:26	00:13.50

Table 3: Comparison of striker test times over rUNSWift's past three years

Another measure of the kick's performance was at the 2012 competition. Here rUNSWift scored the most goals out of any team, scoring 62 goals over 8 games. The teams that came first and second overall scored 54 and 55 goals each, leaving them at least 7 goals behind rUNSWift. Although other aspects of the game play no doubt affected the number of goals scored by rUNSWift, the robustness of the kick was certainly a significant factor in the high goal tally.

10 Concluding Discussion on Team Performance

The true test of performance for such a developmental system is running it in the environment it was designed for, seeing if it achieves its objectives and observing how it compares to other competitors. In this case that means on the soccer field under the pressure of competition matches.

In 2012 rUNSWift finished 3rd place overall and 2nd in the Open Challenge, a big improvement on the 2011 results where we were eliminated in the Quarter Finals. rUNSWift was also the top goal scoring team in the SPL, scoring a total of 62 goals over 8 games. This was again a huge improvement on the 2011 results of scoring 23 goals over 6 games. The team also did not score any own goals at the competition, which was a big achievement considering the change to a symmetric field layout. Figure 23 shows the results of the semi finals onwards, with the top 4 teams coming from the University of Texas at Austin (Austin Villa), the University of Bremen (BHuman), the University of Leipzig (Nao-Team HTWK) and the University of New South Wales (rUNSWift).

Semi Finals

S1	**B-Human**	Nao-Team HTWK	2 : 2 (4 : 3 after penalties)
S2	**Austin Villa**	rUNSWift	7 : 6

3rd Place

SF	Nao-Team HTWK	**rUNSWift**	1 : 11

Final

F	B-Human	**Austin Villa**	2 : 4

Figure 23: 2012 Competition Results

References

[Aldebaraan-Robotics, 2012a] Aldebaraan-Robotics. Nao sonar setup, 2012.

[Aldebaraan-Robotics, 2012b] Aldebaraan-Robotics. Nao technical specifications, 2012.

[Anderson, 2012] Peter Anderson. New Methods for Improving Perception in RoboCup SPL. Honours Thesis, The University of New South Wales, 2012.

[Ashar *et al.*, 2010] Jayen Ashar, David Claridge, Brad Hall, Bernhard Hengst, Hung Nguyen, Maurice Pagnucco, Adrian Ratter, Stuart Robinson, Claude Sammut, Benjamin Vance, Brock White, and Yanjin Zhu. RoboCup standard platform league - rUNSWift 2010. In *Australasian Conference on Robotics and Automation*, 2010.

[Chen and Medioni, 1991] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Proceedings of the IEEE International Conference on Robotics and Automation*, 3:2724–2729, 1991.

[Cohen, 2005] Paul R Cohen. If not Turings's test, then what? *AI Magzine*, 26(4):61–67, Winter 2005.

[Hunter, 2012] Youssef Hunter. Humanoid Robot Localisation for the RoboCup Standard Platform League. Undergraduate Honours Thesis, The University of New South Wales, 2012.

[Kalman, 1960] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering Transactions*, 82, 1960.

[Konforti, 2006] Boyana Konforti. More than the sum of its parts. *AI magazine*, 27(4):19–34, January 2006.

[Kurniawan, 2011] Jimmy Kurniawan. Robot Detection using Vision. Undergraduate Honours Thesis, The University of New South Wales, 2011.

[Mahalanobis, 1936] Prasanta Chandra Mahalanobis. On the generalised distance in statistics. In *Proceedings of the National Institute of Sciences of India*, volume 2, pages 49–55, 1936.

[Teh, 2012] Belinda Teh. Dynamic Omnidirectional Kicks on Humanoid Robots. Honours Thesis, The University of New South Wales, 2012.