rUNSWift Team Description

Claude Sammut[1] and Inuka Aamaratunga[1] and Joshua Chan[1] and Willian Chan[1] and Andrew Chung[1] and Stefan Immaraj[1] and Martin Velikov[1] and Eric Zhou[1]

[1] School of Computer Science & Engineering, The University of New South Wales, Sydney 2052, Australia
c.sammut@unsw.edu.au

Abstract. rUNSWift has a long history in RoboCup, having won the SPL championship several times. We are transitioning to the new Humanoid League with the Booster K1. Our software implementation combines pre-exiting walk and vision models with our own code for localisation and behaviours as well as new locomotion models. A novel contribution is the development and use of the TeleoR robot programming language. Future development will replace all pre-existing code with our own, including new reinforcement learning methods.

Keywords: Humanoid Robot, Reinforcement Learning, Teleo-reactive Programming.

1 History and Contribution

UNSW has participated in the standard platform league every year since 1999 (when it was the SONY Aibo league). We have won the championship five times (2000, 2001, 2003) with the Aibo and in 2014, 2015 with the Aldebaran NAO robot.

One of our most significant contributions to the league has been the development of the walking engine that became the basis of most of the SPL teams. In 2003, we become the first team in RoboCup to use reinforcement learning to refine a robot's walk [1]. RL was again used on the NAO [2]. The speed and agility of the walking engine was a major contributor to our wins in 2014 and 2015. Subsequently, B-Human took our code and improved on it, so that is the code that many of the other teams have used.

Our work in RoboCup has also led to publications in robot vision, localisation and machine learning [3-10].

1.1 Impact on Graduates

A popular course in the Computer Science and Engineering curriculum is "Robot Software Architectures", which is a project-based course where RoboCup related projects are encouraged and used to both educate students and prepare them for membership in the RoboCup team. This has had a profound effect on the students' education as it introduces them to problems when dealing with physical devices with all the uncertainty of the real-world.

We have found that our RoboCup students are highly regarded in industry with graduates going to senior positions in companies like Google, DeepMind, Ubber, Meta, SpaceX, Toyota and Honda Research, as well as creating their own robotics startups, such as Reimagine Robotics.

1.2 Local Community

rUNSWift has been a strong supporter of RoboCup Junior Australia, hosting finals competitions and demonstrating robots from the major competitions. This and visits to schools have inspired students to engage with STEM subjects and later enter university courses in Computer Science and Mechatronics, sometimes going on to join our teams.

In 2025, we also conducted demonstrations for a UNSW organised symposium on the societal impact of AI and Robotics, which was open to the public and attracted a very large and diverse audience.

2 Lessons and Challenges

Coming from the Standard Platform League, a very clear lesson is that being able to share software makes an enormous contribution to the development of the league. New teams benefit by being able to enter without too much pain and experienced teams can build on each other's successes so that we are not all duplicating effort.

Having a well-constructed platform also means that we can focus on the software, which is where most progress is made. The SPL has evolved sophisticated behaviours and team play because we could devote our efforts to software development.

We to keep these advantages for the new Humanoid League.

2.1 Moving to the Booster K1

With the integration of the SPL and Humanoid leagues, rUNSWift has purchased five Booster K1 robots that were delivered in late December 2025. This is a major change from the NAO robots, but we have had reasonable preparation for the change. In 2025, we made the critical decision to port our NAO code to ROS 2. Although this had its problems, making such complex code work in a new framework, it has given the team a lot of experience with ROS, which will carry over to the Booster. In addition, our team, combined with NUbots, participated in the Beijing Masters Tournament, part of the World Humanoid Robot Games in 2025. Our team learned to program the Booster T1 and with only a week's preparation reached the semi-finals, being defeated by the eventual winner, Tsinghoa, who had much more experience with these robots.

In the limited time before the 2026 competition, it is unlikely that we can port all our code. However, Booster has provided sample code for RoboCup, which is a help, but not a complete solution. At present, their localisation code does not work on the K1, so we will have to rewrite that. Furthermore, locomotion is limited to walking, so we will have to train new kicks and movements. Their game play behaviours are also very limited. It is only for a 3 vs 3 game, and lacks much cooperation between robots, so these are the areas we are focussing on for 2026.

3 Plans for 2026

The Booster sample code for the K1, built on ROS 2 Humble provides:

- a pre-trained walk (no soccer-specific actions)
- vision for ball and landmark detection,
- localisation (broken),
- behaviours using a behaviour tree formalism (limited).

In the long term, we want to use our own code for walking and vision, but this may not be practical in this first year with the K1s. So, our implementation strategy is:

- Use Isaac Gym learning environment to train kicks and integrate them with the existing walk model.
- Keep the existing vision models as they seem reasonably reliable
- Replace Booster's localisation with our own.
- Replace Booster's behaviour tree with our own implementation of teleo-reactive programming.

3.1 Reinforcement Learning

As part of the Robot Software Architectures course in 2025, members of the RoboCup team have begun developing new kicks, using BoosterGym [11] and HTWK-Gym[1] which are reinforcement learning environments built on Nvidia's Issac Gym. This includes a policy proximal learning algorithm that we are currently using. However, it has a very complicated reward function with many parameters, so considerable experimentation is required to achieve a satisfactory behaviour. We continue to use this approach for now, but in the longer term, we will explore our own learning methods that combine qualitative reasoning with reinforcement learning [10] to achieve better data efficiency.

We also have experience controlling the robot's actions from our participation in the Beijing Masters, so we are not entirely reliant on machine learning.

3.2 Localisation

Booster's sample RoboCup code includes a particle filter for localisation. We believe this was ported from the T1 implementation, but we've been unable to make it work so far, so have to revert to our own localisation code.

The rUNSWift localisation is based on a multi-model Kalman filter that was developed for the 2014/2015 competitions and has evolved since then. A challenge will be integrating this with Booster's landmark recognition as our previous vision model was based on field line detection, whereas the Booster model looks for corners and T-junctions.

[1] https://github.com/NaoHTWK/htwk-gym

3.3 Behaviours

Booster has implemented its behaviours using Behaviour Trees, where the behaviours are specified in an XML file. This is not ideal as the behaviours are very simple, basically go to the ball and kick it towards the goal. There is very little team coordination. In addition, the XML formalism is hard to read and very cumbersome to write.

UNSW also participates in the RoboCup@Home competition. As part of that, we have collaborated with Prof. Keith Clark at Imperial College London and Dr. Peter Robinson at the University of Queensland, implementing the @Home behaviours using the TeleoR language[2]. This is also well suited to describing soccer behaviours and we already have an implementation for ROS. With the arrival of the K1s, we are porting TeleoR to the humanoid robots.

4 Implementation Status

- Walk, based on Booster model, is working. Additional actions are currently being trained on Issac Gym.
- Vision, also based on Booster models, is working.
- Localisation using our own code has been written and is being tested.
- Behaviours are being rewritten in TeleoR. The language is currently being ported.

5 Potential Joint Team

We have been in discussion with RMIT RedbackBots and the University of Newcastle NUbots about creating a joint team, if needed. The RMIT team is led by Dr. Tim Wiley, a UNSW graduate and former member of rUNSWift, and we previously partnered with the NUbots, led by Prof. Stephan Chalup, in the Beijing Masters. Apart from the research benefits of collaborating with well-known partners, a joint team would help spread costs. Although rUNSWift has five K1s, shipping all of them will be very expensive, so we would prefer to send a smaller number to Incheon.

Acknowledgments. The rUNSWift code has been built on the legacy of all the previous teams, tracing back to 1999. All the teams have been supported by the School of Computer Science and Engineering. Our teams are also grateful to the Office of the New South Wales Chief Scientist and Engineer for travel grants.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

[2] https://staff.itee.uq.edu.au/pjr/HomePages/QulogFiles/teleor_user_guide.pdf

References

1. B. Hengst, S. B. Pham, D. Ibottson, and C. Sammut. Omnidirectional locomotion for quadruped robots. In RoboCup 2001: Robot Soccer World Cup V, pp. 368–373, Berlin: Springer (2002)
2. J. Ashar, J. Ashmore, B. Hall, S. Harris, B. Hengst, R. Liu, Z. Mei (Jacky), M. Pagnucco, R. Roy, C. Sammut, O. Sushkov, B. Teh, and L. Tsekouras, RoboCup SPL 2014 champion team paper. In RoboCup 2014: Robot World Cup XVIII (R. A. C. Bianchi, H. L. Akin, S. Ramamoorthy, and K. Sugiura, eds.), vol. 8992 of Lecture Notes in Computer Science, pp. 70–81, Springer International Publishing (2015).
3. A. Ratter, D. Claridge, J. Ashar, and B. Hengst. Fast object detection with foveated imaging and virtual saccades on resource limited robots. In AI 2011: Advances in Artificial Intelligence (D. Wang and M. Reynolds, eds.), (Berlin, Heidelberg), pp. 560–569, Springer Berlin Heidelberg (2011).
4. P. Anderson, Y. Yusmanthia, B. Hengst, and A. Sowmya. Robot localisation using natural landmarks. In RoboCup Symposium, Mexico City. (2012).
5. B. Hengst. Reinforcement Learning of Bipedal Lateral Behaviour and Stability Control with Ankle-Roll Activation. In Nature-Inspired Mobile Robotics. pp. 411–418. (2013).
6. K. Clark, B. Hengst, M. Pagnucco, D. Rajaratnam, P. Robinson, C. Sammut, and M. Thielscher. A framework for integrating symbolic and sub-symbolic representations. In 25th International Joint Conference on Artificial Intelligence IJCAI-16, pp. 2486–2492, New York, New York, USA, (2016).
7. W. Lu and C. Sammut. Object recognition with class conditional gaussian mixture model - a statistical learning approach. In RoboCup Symposium (2022).
8. W. Lu and C. Sammut. D-flow: A real time spatial temporal model for target area segmentation. In IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids). Okinawa, Japan. (2022).
9. W. Lu, Q. Sheh, L. Li, and C. Sammut. Object recognition with class conditional gaussian mixture model - a statistical learning approach. In Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 13561 LNAI, pp. 3–13. (2023).
10. T. Wiley and C. Sammut, Data efficient online learning of robot behaviours via qualitative planning and reinforcement learning, Robotics and Autonomous Systems, vol. 194, Dec (2025).
11. Y. Wang et al., BoosterGym: A Scalable and Unified Benchmark for Humanoid Robot Reinforcement Learning. arXiv preprint arXiv:2506.15132v1. (2025). https://arxiv.org/abs/2506.15132v1