

Team Description Paper whIRLwind Amsterdam

Mark Honkoop, Julia de Vries, and Gijs de Jong

University of Amsterdam, Science Park 900, Netherlands

`info@whirlwind.team`

`http://www.whirlwind.team`

Abstract. whIRLwind is a student team at the University of Amsterdam that participates in the RoboCup Humanoid Soccer League. The team formerly competed in the Standard Platform League (SPL) as the Dutch Nao Team. This document describes the software projects that the team has worked on over the past year, including RL behaviors and several RL locomotion projects, as well as the team’s impact on both the RoboCup and the university’s community.

Keywords: Team description, whIRLwind, Reinforcement Learning

1 Team Information

Team whIRLwind Amsterdam, formerly known as The Dutch Nao Team, from the Netherlands comprises of 15 active members and a staff member from the University of Amsterdam. The team consists of students pursuing degrees in both Bachelor’s and Master’s programs in Artificial Intelligence and Computer Science. Starting this year, the team will be participating in the middle robot division of the Humanoid Soccer League.

2 Own Contribution

Although we plan to attend RoboCup 2026 with Booster K1 robots, we have not yet received the robots at the time of writing. This, combined with the fact that our team was previously participating in the SPL and not in the humanoid league, means that this year’s projects are still under heavy development.

Nevertheless, whIRLwind will be making a lot of significant advancements this year that will hopefully be ready to be used at RoboCup 2026. These advancements include the development of a new framework called **maelstrom**. We are investing significant time and resources in Reinforcement Learning (RL). We currently use RL for our kick-motion, locomotion, goalkeeper and behaviors. An elaborate overview of all the projects conducted in 2025 can be found in the 2025 team report [6]

2.1 Framework

Our team has been participating in the RoboCup events since 2024 with `yggdrasil`, our framework fully built in the Rust programming language. All functionalities in this framework have been developed by the Dutch Nao Team, and the framework does not use any other team’s code. Because our team is switching from the NAO robots to Booster K1’s, there is suddenly a lot more compute power available to us, which is why we decided to work on a new framework called `maelstrom`.

2.2 RL behaviours

The past two years, our team has worked on training RL behaviors in our own custom 2D physics framework, DNT-RL [6]. Building on those efforts, our team will be working on more complex RL behaviors this year. Just as our previous framework, our new framework will be using a modular behavior system that combines both algorithmic and RL behaviors.

One of the RL behaviors that was successfully tested and deployed at the RoboCup 2025 is the **search behavior**. Here the robot’s task is to locate the ball using a limited field of view. The environment models a vision cone and a discretised field, encouraging systematic exploration and strategic scanning rather than random wandering. Success is defined in perceptual terms (bringing the ball into view), so the learned policy must coordinate motion and orientation to probe unseen regions of the field efficiently while staying within safe bounds.

2.3 RL locomotion

We have created a simulated training environment for the K1 robot and have been working on several RL locomotion tasks, such as walking, running, passing/kicking and goalkeeping. Within these tasks, we have implemented and are experimenting with the use of imitation learning techniques, such as Adversarial Motion Priors (AMP) [13].

RL kick As of now we have a prototype policy for ball passing/kicking, where the robot is able to kick the ball to a specific target. For this task we have been experimenting with leveraging knowledge about symmetry through the use of data augmentation and mirror loss as described by [10], as well as through using equivariant neural networks [14].

RL goal keeper Currently we have a working prototype of the goal keeper using hand-crafted ball observations. However, we are currently also exploring the feasibility of a fully end-to-end policy for the goalkeeper, using raw data from the camera as input. The goal is to come up with an efficient world model that the policy can update. Currently, we are working on a technique inspired by [18] combined with AMP [13] using motion capture data we recorded in our lab.

Walking engine We currently use a fully RL-based locomotion policy that conditions on a commanded target velocity provided as part of the observation. The policy is trained entirely in simulation using RSL-RL [17] and PPO [16], and we apply symmetry augmentations and AMP [13] (using motion capture data we recorded in our lab) to create a stable and stylized gait. The policy is trained using a structured curriculum, progressively increasing the requested velocities. To transfer the policy from sim to real, we apply several techniques, ranging from domain and dynamics randomization to DAgger-style [15] behavior cloning, to create a policy that generalizes to the real world.

3 Impact

The active participation of whIRLwind (called the Dutch Nao Team at the time) in RoboCup over the past years has had significant impact on the team itself, the University of Amsterdam, and the SPL. Over the years, the team has supported several theses and projects that led to publications [9,12,5,3,2,11,1] and projects that led to publications on a large variety of topics [8,4,7]. Additionally, whIRLwind is part of the Lab42 community¹, which is the center for Digital Innovation and AI of the University of Amsterdam. Finally, the Dutch Nao Team was the only team in the Challenge Shield of the SPL that developed its own framework, something we believe provided a significant input to the SPL and encouraged new development and research.

Aside from the events at the University of Amsterdam that whIRLwind participates in, whIRLwind organizes and is part of a lot of events for the entire community. The team has a collaboration with Startup Village Amsterdam². In practice, this means that delegations from all over the world come and visit the Intelligent Robotics Lab and the team gives presentations and demonstrations on the work we do as a team. Furthermore, we give programming workshops for high-school students, both in our robotics lab and at events like the Career Day³. Lastly, we participate in events that are free for the entire community, such as 24UurOost⁴, an event organized by the city of Amsterdam, where we give presentations about the work that we do for the people of Amsterdam.

References

1. Bernardy, A.M.: Introducing object detection to ekf slam on a nao-robot football field
2. gezegd Deprez, H.L.: Enhancing simulation images with gans. Bachelor thesis, Universiteit van Amsterdam (2020), https://staff.fnwi.uva.nl/a.visser/education/bachelorAI/thesis_hidde_lekanne_deprez.pdf

¹ <https://lab42.uva.nl/>

² <https://www.startupvillage.nl/>

³ <https://www.techportal.nl/career-day>

⁴ <https://www.iamsterdam.com/uit/24uur-amsterdam/oost>

3. Garritsen, T.: Using the extended information filter for localization of humanoid robots on a soccer field. Bachelor's thesis, Universiteit van Amsterdam (6 2018)
4. van Heusden, R.: Making a robot stop a penalty. In: Proceedings of the 30th Belgian-Netherlands Conference on Artificial Intelligence (BNAIC). pp. 89–90 (11 2018)
5. van Heusden, R.: Making a robot stop a penalty – using q learning and transfer learning. Bachelor's thesis, Universiteit van Amsterdam (6 2018)
6. Jansen, V., Jansen, A., Honkoop, M., Blaauboer, J., Meijer, M., Oud, M., van der Veen, R., Nagelhout, F., Ruiter, H., Gunnewiek, F.K., Verkuijl, K., Catarrinho, D.X., de Jong, G.: Dutch nao team - team report (2026)
7. de Jong, G., Eshuijs, L., Visser, A.: Learning to walk with a soft actor-critic approach (11 2023)
8. Lagrand, C., van der Meer, M.: The roasted tomato challenge. *Amsterdam Science* **05**, 4 (4 2017)
9. Lagrand, C.G.: Learning a robot to score a penalty – minimal reward reinforcement learning. Bachelor's thesis, Universiteit van Amsterdam (6 2017)
10. Mittal, M., Rudin, N., Klemm, V., Allshire, A., Hutter, M.: Symmetry considerations for learning task symmetric robot policies. In: 2024 IEEE International Conference on Robotics and Automation (ICRA). pp. 7433–7439. IEEE (2024)
11. Monté, X.: Neural factorization of shape and reflectance of a football under an unknown illumination. Bachelor thesis, Universiteit van Amsterdam (2023), https://staff.fnwi.uva.nl/a.visser/education/bachelorAI/Monte_NERF_of_a_football_under_unknown_illumination.pdf
12. Negrijn, S.: Exploiting Symmetries to Relocalise in RoboCup Soccer. Master's thesis, Universiteit van Amsterdam (12 2017)
13. Peng, X.B., Ma, Z., Abbeel, P., Levine, S., Kanazawa, A.: Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (ToG)* **40**(4), 1–20 (2021)
14. Van der Pol, E., Worrall, D., van Hoof, H., Oliehoek, F., Welling, M.: Mdp homomorphic networks: Group symmetries in reinforcement learning. *Advances in Neural Information Processing Systems* **33**, 4199–4210 (2020)
15. Ross, S., Gordon, G., Bagnell, D.: A reduction of imitation learning and structured prediction to no-regret online learning. In: Proceedings of the fourteenth international conference on artificial intelligence and statistics. pp. 627–635. JMLR Workshop and Conference Proceedings (2011)
16. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
17. Schwarke, C., Mittal, M., Rudin, N., Hoeller, D., Hutter, M.: Rsl-rl: A learning library for robotics research. arXiv preprint arXiv:2509.10771 (2025)
18. Zhang, C., Klemm, V., Yang, F., Hutter, M.: Ame-2: Agile and generalized legged locomotion via attention-based neural map encoding. arXiv preprint arXiv:2601.08485 (2026)