Badger Bots Team Description Paper

Adam Labiosa, Will Cong, Josiah Hanna

University of Wisconsin–Madison

1 Introduction

We are team Badger Bots from the University of Wisconsin–Madison, Madison, Wisconsin, USA. Our team consists of graduate and undergraduate students from the Department of Computer Sciences. Over the past three years, we have competed in the RoboCup Standard Platform League (SPL) under the team names Badger Bots and (as a part of) WisTex United. As WisTex United, we formed a joint team with the University of Texas at Austin (UT Austin Villa).

Our recent competition results include 3$^{\text{rd}}$ place in the Champions Cup 2025, 1$^{\text{st}}$ place in the Challenge Shield 2024, and 3$^{\text{rd}}$ place in the Challenge Shield 2023. Our primary research focus is using reinforcement learning (RL) to create intelligent multi-agent behaviors on physical humanoid robots. Toward this goal, we have developed and deployed RL-trained behaviors on NAO robots for the past three years.

2 Lessons Learned from RoboCup Participation

Participation in previous RoboCup competitions has provided lessons regarding simulation, physical robot testing, and research vs. engineering balance. Since our work focuses on deploying RL-trained policies on real robots, many of these lessons stem from sim2real challenges.

High-fidelity simulation as a development tool. During competitions, the time available to modify, deploy, and test code is extremely limited. We have learned that high-fidelity simulation is essential for reducing iteration time and catching errors. In our workflow, we first build our code in simulation to identify basic issues such as compilation errors. Second, we test behaviors in a lightweight simulation without computationally expensive vision and localization to verify correctness on a single robot. Finally, we evaluate behaviors in scenario-specific simulations (e.g., near-goal situations). This approach enables fast iteration within the competition's time constraints.

Simulation has also allowed us to safely evaluate aggressive or unconventional behaviors without risking hardware damage. For example, in our first year we relied on a shielding mechanism to limit overly aggressive behavior. Simulation enabled us to tune and visualize these mechanisms before deployment, avoiding potentially damaging robot collisions.

Need for physical testing. Despite its importance, we have learned not to over-rely on simulation. Policies that perform well in simulation can fail on real robots due to the sim2real gap. In earlier competitions, insufficient on-robot testing forced us to make large between-game and sometimes in-game adjustments, which led to suboptimal performance. As a result, we prioritize deploying policies on physical robots early and validating them under real-game conditions.

We have also learned the importance of explicitly testing situations that may appear to be edge cases but are critical in-game situations. For example, ball search behaviors, near-goal interactions, and changes in vision conditions. Failures in such situations can significantly reduce scoring opportunities and limit performance.

Engineering vs. research balance. RoboCup requires a careful balance between novel research and engineering-focused system integration. Throughout most of the year, our efforts focus on research, such as new behavior learning environments. However as the competition approaches, we have learned to shift toward robustness, reliability, and effective on-robot performance.

Although this transition can appear to move away from traditional research, we have learned that the competition often produces equally valuable and publishable results. The competition environment exposes full-system challenges that are difficult to find in isolated experiments, and addressing these challenges has lead us to insights in architecture design, sim2real transfer, and other learning methodologies.

3 Major Challenges

For the RoboCup 2026 competition, our team is addressing three major challenges: core robot functionality, reinforcement learning for soccer, and high-fidelity simulation.

Core robot functionality. We are developing a modular robot software stack that supports reliable RL-based soccer play. This includes integrating the Booster Robotics RoboCup demo framework [1] with localization modules from B-Human [6], while adding support for RL-based policy inference. Key challenges include full-system reliability, robot safety and correct in-game behavior.

Reinforcement learning for soccer. Our teams long-term objective is to enable intelligent on-robot behavior through RL at both the locomotion and decision-making levels. For locomotion, we are developing RL-trained walking and soccer-specific skills such as kicking and dribbling. We are extending MuJoCo Playground [8] to support the Booster K1 humanoid and are validating learned policies in high-fidelity RoboCup simulations and on physical robots. We are focusing on the major challenge of bridging the sim2real gap to enable effective on-robot locomotion.

At the behavior level, we are training high-level soccer policies using RL with MJX. This includes multi-agent training, RL-based self-play, and integration with learned locomotion skills. Our goal is to produce effective multi-agent soccer behaviors that operate within the RoboCup rules and integrate effectively with our locomotion policies.

High-fidelity simulation. We are also developing a full-stack, high-fidelity simulation environment by extending the RoboCup Simulation League simulation to support inference from our software stack on a simulated Booster K1 robot. This simulation will enable fast, multi-robot and full-field experimentation and will support both development and testing for our team and potentially other SPL teams.

4 Major Changes and Implementation Status

For RoboCup 2026, we anticipate several major changes to both our hardware and software systems. For hardware, we are currently working with the University of Wisconsin–Madison to acquire Booster K1 robots and expect to have them before the competition for testing and evaluation.

On the software side, we aim to deploy a complete robot architecture on the K1 platform. This includes the ROS 2-based Booster Robotics framework [1], B-Human localization [6], and our RL-based locomotion and behavior control policies. Currently, the Booster framework is connected to our high-fidelity simulation setup, localization improvement is ongoing, and our RL-trained locomotion policies are in the process of being transferred to the high-fidelity testing simulation.

5 Team and Research Impact

Through RoboCup participation, our team has contributed to advancing learning-based methods within the SPL and to outreach within our local community. We have worked to spread the use of RL for behavior learning in the SPL, releasing our high-level behavior training environments which allowed multiple other teams to explore RL approaches for their own systems. As we move to a new hardware platform, we aim to continue leading and helping enable research on RL for humanoid soccer.

Beyond the SPL, our team actively engages with the Madison community through robotics demonstrations for middle and high school students and by providing undergraduate research opportunities. Humanoid soccer serves as an accessible platform for introducing students to robotics and AI, and our team has hosted multiple demonstrations through the Computer Sciences department for local students. As well, participation with our team enables many graduate and undergraduate students to gain hands-on research experience by contributing to a large-scale robotics codebase and conducting experimental research with real robots.

6 Publishable Work

RoboCup has served as both a benchmark and a motivation for our research. Following the 2024 competition, we published work describing the RL-based system used to win the Challenge Shield [5]. Additional publications inspired by RoboCup include work on multi-agent RL on the NAOv6 robots [4], vision data augmentation [7], RL data augmentation for robust behavior learning [2], and learning abstractions to improve simulator realism [3].

Looking ahead, we anticipate several publishable directions from our 2026 efforts, including a system-level description of RL-based control on the Booster K1 platform, high-level behavior learning with MJX and self-play, and methods for fine-tuning locomotion using data collected during behavior learning. We expect RoboCup to continue to serve as inspiration for impactful RL and robotics research.

7 Conclusion

Our team, Badger Bots, has been an active and successful participant in the RoboCup SPL for multiple years. Our work emphasizes reinforcement learning for humanoid soccer and we have demonstrated competitive success and and have a record of publishable research. We believe our efforts toward RoboCup 2026 will result in both competitive performance on the field and meaningful contributions to humanoid robotics research.

References

1. BoosterRobotics: BoosterRobotics/robocup_demo (Jan 2026), https://github.com/BoosterRobotics/robocup$_demo, original - date : 2024 - 11 - 25T11 : 22 : 59Z$
2. Corrado, N.E., Qu, Y., Balis, J.U., Labiosa, A., Hanna, J.P.: Guided Data Augmentation for Offline Reinforcement Learning and Imitation Learning (Aug 2024). https://doi.org/10.48550/arXiv.2310.18247, http://arxiv.org/abs/2310.18247, arXiv:2310.18247 [cs]
3. Deng, Y., Hanna, J.: Abstract Sim2Real through Approximate Information States
4. Labiosa, A., Hanna, J.P.: Multi-Robot Collaboration through Reinforcement Learning and Abstract Simulation (Mar 2025). https://doi.org/10.48550/arXiv.2503.05092, http://arxiv.org/abs/2503.05092, arXiv:2503.05092 [cs]
5. Labiosa, A., Wang, Z., Agarwal, S., Cong, W., Hemkumar, G., Harish, A.N., Hong, B., Kelle, J., Li, C., Li, Y., Shao, Z., Stone, P., Hanna, J.P.: Reinforcement Learning Within the Classical Robotics Stack: A Case Study in Robot Soccer (Mar 2025). https://doi.org/10.48550/arXiv.2412.09417, http://arxiv.org/abs/2412.09417, arXiv:2412.09417 [cs]
6. Röfer, T., Laue, T., Böse, F., Cihasev, A., Dehmlow, K., Gittner, L., Hasselbring, A., Lienhoop, J., Lührsen, A., Meinken, Y., Monnerjahn, L.M., Oppermann, M., Plecher, L., Reichenberg, P., Sablotny, R., Schiefelbein, L., Schreiber, S., Struckhoff, T.: B-Human code release documentation 2025 (2025), only available online: https://docs.b-human.de/coderelease2025/
7. Shao, Z., Hanna, J.P.: WeRef: An Open-source and Extensible Dataset for Referee Gesture Recognition in RoboCup. In: RoboCup-2025: Robot Soccer World Cup XXVIII (Jul 2025)
8. Zakka, K., Tabanpour, B., Liao, Q., Haiderbhai, M., Holt, S., Luo, J.Y., Allshire, A., Frey, E., Sreenath, K., Kahrs, L.A., Sferrazza, C., Tassa, Y., Abbeel, P.: MuJoCo Playground (Feb 2025). https://doi.org/10.48550/arXiv.2502.08844, http://arxiv.org/abs/2502.08844, arXiv:2502.08844 [cs]