

Tsinghua Hephaestus 2026 Team Description Paper

Changsheng Luo, Penghui Chen, Yushi Wang, and Mingguo Zhao*

Department of Automation, Tsinghua University, Beijing, China

Abstract. Tsinghua Hephaestus, the RoboCup 2025 Adult Size champion, presents its research directions for RoboCup 2026. Building on our learning-based framework that demonstrated effective integration of vision-driven control, locomotion, and fall recovery, we identified key limitations in skill diversity and multi-agent coordination during competitive play. For RoboCup 2026, our research focuses on expanding the behavioral repertoire of humanoid soccer robots and developing coordination mechanisms through learning-based approaches. We investigate methods to train diverse soccer skills within unified frameworks and explore the integration of human demonstration data with multi-agent reinforcement learning. Our contributions include open-source simulation infrastructure and validation of end-to-end learning paradigms in competitive humanoid robotics.

Keywords: Humanoid Robot · Robot Soccer · Deep Reinforcement Learning · Multi-Agent Learning

1 Introduction

Since 2004, Tsinghua Hephaestus has competed in the RoboCup humanoid League, advancing from model-based control methods to learning-based techniques. In RoboCup 2025, the team won the Adult Size championship [2].

Our humanoid robots use a learning-based system that closely links vision (robotic visual processing) with adversarial motion priors (movement patterns learned from human demonstration) [4], enabling direct interaction between sensory inputs and the robot’s motors [7]. Key features of our robot include omnidirectional locomotion (ability to move in any direction), curriculum-based fall recovery learning (stepwise training to get up after a fall) [1], and proprioceptive localization (the robot’s sense of its own body position).

Building on the RoboCup 2025 experience, this paper presents lessons learned, outlines challenges to be addressed, and specifies the main objectives and key strategic steps planned for RoboCup 2026.

* Corresponding author, mgzhao@mail.tsinghua.edu.cn

2 Lessons Learned from RoboCup 2025

2.1 Strengths of Learning-Based Approaches

The championship result demonstrated that learning-based methods are highly effective in competitive humanoid robotics. The locomotion subsystem, trained using policy gradient algorithms [5], demonstrated strong mobility and robustness to physical disturbances. The vision-driven control system [7] used motion priors from demonstrations [4] to adapt gameplay without hand-engineered rules. The fall recovery system [1] used curriculum learning to handle many falling situations. Proprioceptive models enabled robots to determine their position using learned representations from internal sensors, while simulation-based training [3] made skill transfer to real robots more efficient.

2.2 Identified Limitations

Although we have achieved success in a single robot’s kicking skills, limitations remain, such as insufficient behavioral diversity and weak collaboration among multiple robots. The system only supports a narrow range of skills for complex matches. Additionally, limited observation contexts and training hinder the development of diverse strategies and coordinated team behaviors.

3 Major Challenges for RoboCup 2026

Advancing humanoid soccer systems requires tackling foundational research challenges, including increasing behavioral diversity in reinforcement learning while ensuring training stability, facilitating robust multi-agent coordination, and transferring high-level strategic knowledge to robots via demonstration-based learning. These interdependent challenges drive ongoing research in the field.

4 Plans for RoboCup 2026

Our development focuses on extending learning-based methods to address identified limitations by expanding behavioral capabilities and maintaining robustness. We are developing training approaches for diverse soccer skills, leveraging both simulation and demonstration data, and exploring coordination mechanisms for multi-agent behaviors.

5 Current Implementation Status

We have established foundational infrastructure based on our existing learning systems. Initial experiments with training methodologies and coordination approaches are in progress. Core capabilities from RoboCup 2025 remain operational and serve as the foundation for ongoing development.

6 Impact and Contributions

6.1 Impact on the Humanoid Soccer League

We contribute to the community through open-source resources and technical validation. The BoosterGym simulation framework [6] provides infrastructure for learning-based research on humanoid robots. Our championship performance validated the effectiveness of learning-based approaches in the AdultSize competition. We propose a hierarchical technical framework that provides a systematic roadmap for developing and evaluating capabilities in humanoid soccer systems.

6.2 Research Contributions

Our contributions include fall recovery through curriculum learning [1], proprioceptive localization, and vision-driven motor control [7], extending motion prior frameworks [4]. Ongoing research focuses on expanding behavioral capabilities, coordination, training diverse soccer skills, and transferring strategic knowledge.

Acknowledgments. This research was supported by STI 2030-Major Projects 2021ZD0201402. We thank Booster Robotics for providing the robot platform, experimental environment, and technical support.

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

References

1. Chen, P., Wang, Y., Luo, C., Cai, W., Zhao, M.: Hifar: Multi-stage curriculum learning for high-dynamics humanoid fall recovery (2025), <https://arxiv.org/abs/2502.20061>
2. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E.: Robocup: The robot world cup initiative. In: Proceedings of the first international conference on Autonomous agents. pp. 340–347 (1997)
3. Makovychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., State, G.: Isaac gym: High performance gpu-based physics simulation for robot learning (2021), <https://arxiv.org/abs/2108.10470>
4. 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)
5. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017)
6. Wang, Y., Chen, P., Han, X., Wu, F., Zhao, M.: Booster gym: An end-to-end reinforcement learning framework for humanoid robot locomotion (2025), <https://arxiv.org/abs/2506.15132>
7. Wang, Y., Luo, C., Chen, P., Liu, J., Sun, W., Guo, T., Yang, K., Hu, B., Zhang, Y., Zhao, M.: Learning vision-driven reactive soccer skills for humanoid robots (2025), <https://arxiv.org/abs/2511.03996>