WHU United Team Extended Abstract

Mingyu Hu, Shangshang Yang, Jiao Zhan, Zihao Mao, Chunyu Li, Jinchun Liu, Jingyang Li, Xiefeng Wu, Meng Mao, Wei Li, Xiangkun Li, Yuwei Tu, Yongxiang Nie, Zhicheng Hu, Yang Li, Jingyi Deng, Yinji Deng, Xiaojing Liu, Yangwen Zhang, Yuhao Guo, and Xinyi Cao

Wuhan University, Wuhan, China
mingyuhu@whu.edu.cn

Abstract. WHU United is applying to participate in RoboCup 2026 in the Humanoid AdultSize (Large Size) League with the Unitree G1 platform. Team members bring prior RoboCup experience from other leagues, including 3D Simulation and the Standard Platform League. Building on this background, we aim to develop a top-down hierarchical decision-making and control scheme with reinforcement learning on the Unitree G1, establishing a strong foundation for competitive performance in RoboCup 2026.

Keywords: RoboCup · Hierarchical control · Reinforcement learning

1 Introduction

1.1 Team and platform overview

WHU United will make its debut in the RoboCup 2026 Humanoid AdultSize League. Based on the Unitree G1 platform, we plan to adopt a top-down hierarchical decision-making and control architecture, systematically integrating reinforcement learning into this framework. We will incorporate proven design concepts from B-Human and Booster-T1, along with our accumulated experience in multi-agent collaboration and visual perception from the 3D Simulation League, to build this competition system. In our first year, we are committed not only to achieving robust and safe operation of the full system throughout matches without crashes but also to progressively validating and deploying our existing technical solutions in key modules such as state estimation, disturbance-resilient whole-body control, and hierarchical multi-agent collaboration. This approach aims to establish a competitive on-field performance foundation.

2 Lessons

Based on our prior RoboCup experience and early development on the Unitree G1 platform, we summarize two key lessons for competing in the Humanoid AdultSize League.

2.1 Low-level control policy is the foundation

A strong low-level control policy plays the most important role in achieving consistent performance on the field. From our experience in the 3D Simulation League, we observed that teams with fast, dynamic, and accurate kicking behaviors can gain decisive advantages. Therefore, our primary goal is to improve the execution speed and accuracy of G1's core skills. In addition, we implement an omnidirectional get-up behavior to reduce recovery time after falls, allowing the robot to quickly return to its assigned role and maintain team structure during matches.

2.2 Multi-agent decision making improves task success rate

In RoboCup soccer, when teams have similar locomotion capabilities, a large portion of the game becomes contention for ball possession. A practical solution in such situations is not always shooting when the goal is blocked, but passing the ball to an appropriate teammate to create better scoring opportunities. Motivated by this, we try to port our passing strategy developed in the 3D Simulation League into a real-world, G1-based multi-agent decision-making system, including coordinated passing and role assignment, to improve overall task success rate and team efficiency.

3 Major Problem and Plans for RoboCup 2026

3.1 Closed-Loop Humanoid Perception and Locomotion

Visual Grounding of Dynamic Objects During Occlusion. Visual grounding in football presents several challenges. Target tracking with onboard cameras is often hindered by motion blur and ghosting artifacts due to continuous motion. Meanwhile, observation often suffers from occlusion, leading to intermittent target disappearance and ambiguous visual cues. Addressing these issues is crucial for establishing a robust and precise perception mechanism that ensures reliable decision-making and control. By overcoming these challenges, we aim to develop a refined and effective perception and transmission framework that enables enhanced performance in dynamic and occluded environments.
State Estimation Under Impact and Vibration Noise. State estimation in robot soccer scenarios must overcome the visual motion blur and sensor noise induced by high-frequency mechanical vibrations, which severely degrade the precision of field feature extraction; Additionally, the instantaneous acceleration saturation and failure of kinematic constraints caused by physical collisions, leading to nonlinear in the Inertial Measurement Unit (IMU). We plan to propose a multi-sensor fusion framework based on the Invariant Extended Kalman Filter (InEKF), providing a robust and reliable foundation for advanced decision-making and dynamic maneuvering in the 2026 RoboCup competition.
Motion Generation Against Dynamic External Perturbations. Humanoid soccer involves highly dynamic, contact-rich interactions, with robots facing unpredictable perturbations such as ball impacts, collisions, sudden ground changes,

and friction variations. Our plan is to propose a framework which allows the robot to execute dynamic soccer behaviors with stability, adaptability, and recovery, providing a robust foundation for higher-level decision-making and multi-agent coordination in RoboCup 2026.

3.2 Multi-Agent Collaboration with Hierarchical Coordination

Existing general-purpose multi-agent collaboration solutions face fundamental challenges in dynamic, adversarial environments. Centralized decision-making is constrained by communication delays and single-point bottlenecks, making it difficult to meet real-time requirements, while fully decentralized approaches, due to the lack of global coordination, often leading to conflicts and inefficient execution. Therefore, how to achieve fast and accurate multi-robot collaborative strategy optimization within limited time has become a core challenge in the field of robot adversarial systems.

4 Current Status

For Visual Grounding, we introduce a detection-oriented motion blur data augmentation strategy to simulates motion-induced imaging degradation during training. Furthermore, leveraging our previous work [5], we adopt a perspective-aware instance segmentation approach which captures the complete shape and semantic information of occluded targets.

For State Estimation, we propose a state estimation framework based on multi-sensor fusion, integrating data from an Inertial Measurement Unit (IMU), legged kinematics odometry, and visual cameras. By employing robust estimation theories such as the Invariant Extended Kalman Filter (InEKF) [1], we achieve real-time computation of robot State and sensor biases. This method exhibits significant robustness in current environments, effectively suppressing non-Gaussian interference resulting from high-frequency mechanical vibrations and instantaneous physical impacts.

For Motion Generation, we develop a disturbance-resilient whole-body control framework that combines a model-based WBC approach with learning-based adaptation. The WBC jointly manages CoM stabilization, contact forces, swing-leg motion, and upper-body posture, prioritizing balance over strict trajectory tracking [3]. Disturbance-aware motion primitives enhance soccer-specific skills. During kicking, impact robustness and accuracy are balanced through coordinated limb motion. During locomotion and defense, low-level stepping strategies replan foot placement, timing, and orientation when balance margins are exceeded, following advances in agile humanoid control on hardware [2].

For Multi-Agent Collaboration, we propose a hierarchical decision-making architecture designed for complex tasks. This architecture follows a two-tier "strategic layer–execution layer" structure. At the execution layer, each robot gathers real-time local environmental data, such as teammate and opponent positions and ball status. The data is preprocessed into a concise, graph-based abstraction, where nodes represent agents, and edges depict spatial relationships. This abstracted information is uploaded to the strategic layer [4]. The strategic layer aggregates the data and uses a policy network, jointly trained with reinforcement and imitation learning, to generate a globally optimized strategy. It then issues abstract instructions in latent space to the execution layer. Upon receiving these instructions, the execution layer combines them with real-time observations to generate specific action sequences, which are ultimately executed.

References

1. Barrau, A., Bonnabel, S.: The invariant extended kalman filter as a stable observer. IEEE Transactions on Automatic Control **62**(4), 1797–1812 (2017)
2. He, T., Gao, J., Xiao, W., Zhang, Y., Wang, Z., Wang, J., Luo, Z., He, G., et al.: Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. arXiv preprint arXiv:2502.01143 (2025)
3. Liao, Q., Truong, T.E., Huang, X., Gao, Y., Tevet, G., Sreenath, K., Liu, C.K.: Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion. arXiv preprint arXiv:2508.08241 (2025)
4. Marzi, T., Alippi, C., Cini, A.: Hierarchical message-passing policies for multi-agent reinforcement learning (2025), https://arxiv.org/abs/2507.23604
5. Zhan, J., Luo, Y., Guo, C., Wu, Y., Yang, B., Wang, J., Liu, J.: Amodal instance segmentation with dual guidance from contextual and shape priors. Applied Soft Computing **169**, 112602 (2025)