UT Austin Villa
2026 Large Robot Division
Extended Abstract

Dongmyeong Lee, Zifan Xu, Nathan Tsoi, Hao Fu, Junhong Xu, Jiaheng Hu, Myoungkyu Seo, Joydeep Biswas, and Peter Stone

The University of Texas at Austin, Austin TX 78712, USA
`domlee,zfxu@utexas.edu`

Abstract. UT Austin Villa is a returning participant in the RoboCup Humanoid League large size competition, building on our initial participation last year. We introduce our team, summarize our current hardware and software infrastructures, and present our short-term development plans to further improve robot agility and perception. We also outline several longer-term research directions—including whole-body control, learning from human soccer videos, and multi-agent coordination—that we aim to pursue using RoboCup as a testbed.

1 Introduction

UT Austin Villa is a returning team in the RoboCup Humanoid League Adult-Size competition, building upon our initial participation last year. The team seeks to advance humanoid robotics through agile motion control, robust sensorimotor skills, and multi-agent strategy, grounded in established robotics and AI research at The University of Texas at Austin [5, 1–4, 6].

We briefly describe our current hardware platform, software capabilities, and development plans for the upcoming competition cycle. These efforts target continued improvements in locomotion, perception, and team coordination informed by lessons learned from our previous participation. Beyond the competition, we aim to leverage RoboCup as a platform for research on whole-body control, visual imitation learning, and multi-agent strategies.

2 Hardware and Software

Our primary hardware platform, *Booster T1*, is a 1.18-meter-tall humanoid robot with 23 degrees of freedom and omnidirectional walking capability. It is equipped with a Stereolabs ZED2i stereo camera for perception and localization, as well as a 9-axis IMU for whole-body control[1]. We currently possess two of these robots, enabling us to explore multi-robot strategies.

Our current software framework is a combination of C++ and Python, with data streaming handled through ROS2. We use the Booster Robotics SDK for

[1] Although the onboard IMU includes a magnetometer, it is currently disabled and not used for odometry.

low-level motor control and for reading proprioceptive data, and we additionally deploy an RL-trained get-up policy provided with the platform. On top of this interface, we run our own locomotion and kicking policies, along with our own odometry estimation. We adapted RoboCup Demo and integrated it into our pipeline. Our system is organized as multiple ROS2 nodes (e.g., perception, localization, and high-level decision making for strike/defense). The perception system uses a fine-tuned YOLOv8 model, and high-level decisions are implemented using behavior trees.

3 Short-Term Plans

Acquiring Diverse Soccer Skills: In our previous work, we have learned advanced striking skill [5]. Our first short-term plan is to leverage human motion priors to learn more diverse set of soccer skills such as goal keeping and saving the ball from the board line. These skills are then sequenced by a behavior tree to execute sophisticated game strategies.
Sensorimotor Policy: Currently, our kicking policy [5] uses ball detections from YOLOv8 as input. In the short term, we plan to use RGB images directly as input to an end-to-end sensorimotor policy to reduce dependence on ball detection accuracy. This approach is expected to enable more robust kicking.

4 Research Interests

Learning Whole-Body Control: We aim to develop controllers that move beyond standard gaits. Through reinforcement learning, we seek agile, whole-body behaviors suited for dynamic soccer scenarios.
Learning from Human Soccer Videos: We will explore using abundant soccer footage to facilitate policy learning. By extracting movement primitives and tactics, we hope to achieve more natural and adaptable robot behaviors.
Multi-Agent Coordination: RoboCup's team-based competitive game setting provides a rich environment for studying multi-agent collaboration, competition, and role allocation. Our goal is to develop teams that adapt to diverse opponents and non-stationary conditions.

References

1. Hu, J., Hendrix, R., Farhadi, A., Kembhavi, A., Martín-Martín, R., Stone, P., Zeng, K.H., Ehsani, K.: Flare: Achieving masterful and adaptive robot policies with large-scale reinforcement learning fine-tuning. arXiv preprint arXiv:2409.16578 (2024)
2. Hu, J., Stone, P., Martín-Martín, R.: Causal policy gradient for whole-body mobile manipulation. RSS (2023)
3. Li, J., Zhu, Y., Xie, Y., Jiang, Z., Seo, M., Pavlakos, G., Zhu, Y.: Okami: Teaching humanoid robots manipulation skills through single video imitation. In: 8th Annual Conference on Robot Learning (2024)
4. Xu, Z., Raj, A.H., Xiao, X., Stone, P.: Dexterous legged locomotion in confined 3d spaces with reinforcement learning. arXiv preprint arXiv:2403.03848 (2024)

5. Xu, Z., Seo, M., Lee, D., Fu, H., Hu, J., Cui, J., Jiang, Y., Wang, Z., Brund, A., Biswas, J., et al.: Learning agile striker skills for humanoid soccer robots from noisy sensory input. arXiv preprint arXiv:2512.06571 (2025)
6. Zhu, Y., Lim, A., Stone, P., Zhu, Y.: Vision-based manipulation from single human video with open-world object graphs. arXiv preprint arXiv:2405.20321 (2024)