CAU Mountain&Sea 2026 Humanoid Soccer Middle Size Extended Abstract

Zhiyuan Zhang, Ye Liu, Xiurong Lv, Ziqian Guo, and Biao Hu[1]

China Agricultural University, Beijing 100083, China
hubiao@cau.edu.cn

Abstract. This paper presents detailed insights into the technical strategies employed to enhance robot performance, with a specific focus on robust locomotion, dynamic action execution, and multi-robot tactical collaboration. Our experiences highlight key innovations in humanoid robotics, including deep reinforcement learning for stable gait generation, adaptive motion control architectures, and distributed coordination frameworks for team strategy. These lessons aim to contribute to the broader field of robotics by showcasing solutions for high-dynamic mobility and cooperative gameplay in competitive humanoid soccer.

1 Lessons We Learned

1. **Collaborative State Estimation Overcomes Individual Perception Limits**: While individual perception is necessary, we learned that relying solely on a single robot's sensors is insufficient in a cluttered environment. Instead of focusing purely on improving individual vision algorithms, we found that *multi-robot information fusion* is the key to robustness. By sharing localization and object data across the team, robots can compensate for individual blind spots or occlusions. Future improvements will focus on "Shared World Models," where the team collectively estimates the ball and opponent positions, allowing robots to make decisions based on a unified team consensus rather than noisy local data.
2. **Modular Architecture Facilitates Scalable Multi-Agent Behavior**: Previously, our decision-making relied on rigid Finite State Machines (FSM). We discovered that FSMs struggle to handle the complexity of multi-agent strategies. To improve scalability, we are transitioning to Behavior Trees (BTs) combined with a dynamic role allocation system. This modular approach allows individual robots to adapt their actions—such as switching from a striker to a supporter role—based on the team's strategic needs rather than just their local state. This flexibility is critical for implementing high-level tactical plays involving passing and coordinated defense.

2 Major Problems

1. **Optimizing Robot Gait for Dynamic Stability and Energy Efficiency**: Efficient locomotion remains a fundamental challenge. The robot

must not only move quickly but also maintain balance during sudden stops, turns, and physical collisions with opponents. Developing an adaptable gait pattern that optimizes stability across different friction conditions is crucial. We are currently addressing the "reality gap" in our reinforcement learning policies to ensure that the agile movements learned in simulation translate effectively to the physical robots. Additionally, we are optimizing the transition phases between walking, running, and kicking to minimize energy consumption and mechanical stress, ensuring the robots can sustain high-performance operation throughout the match.

2. **Real-Time Multi-Agent Coordination and Strategy**: In highly dynamic games, individual intelligence is insufficient; the team must act as a cohesive unit. A major problem we face is the latency and reliability of decision-making when coordinating multiple robots. Robots need to assess their teammates' intent and positions in real-time to execute complex plays, such as passing into space or covering defensive gaps. We are working on resolving conflicts in role assignment (e.g., two robots attempting to interact with the ball simultaneously) by implementing a hierarchical multi-agent planner that prioritizes team utility over individual cost functions.

3 Plans for RoboCup 2026

1. **Advanced Coordination and Communication Framework**: To enhance team synergy, we will implement a low-latency communication protocol that enables robots to share critical game data, including self-localization confidence and intended actions. A dynamic role assignment strategy will be fully deployed, enabling robots to fluidly switch roles based on the game context. For example, a robot nearest to the ball will automatically assume the 'Striker' role, while others adjust their positioning to support or defend. We aim to integrate multi-agent reinforcement learning (MARL) to optimize these strategic role transitions, ensuring the team adapts instantly to opponent behaviors.
2. **Next-Generation Motion Control**: We plan to significantly upgrade our motion engine to support more aggressive and dynamic maneuvers. This includes the development of a "whole-body control" framework that allows the robot to kick the ball while walking without coming to a complete stop, thereby increasing the speed of play. Furthermore, we will refine our push-recovery algorithms, enabling the robot to maintain balance even after significant collisions. By integrating inertial measurement unit (IMU) feedback directly into the gait cycle at a high frequency, we aim to achieve a level of stability that allows for robust gameplay in close-quarters physical engagements.

4 Current Status

We have developed a high-performance humanoid platform, the Booster K1, optimized for agile mobility and cooperative gameplay. The core strength of our current system lies in its robust locomotion engine, which utilizes a hybrid approach of model-based control and reinforcement learning. This allows the robots to approach the ball with optimized walking gaits, execute omnidirectional turns, and perform powerful, stable kicks.

In terms of team intelligence, we have successfully deployed a preliminary multi-robot communication framework. This system allows our robots to maintain awareness of their teammates' positions and execute basic role-switching behaviors, reducing interference between allied robots.

With these advancements in dynamic motion and team coordination, we are well-prepared for the RoboCup 2026 competition, focusing our remaining efforts on refining the fluidity of multi-robot tactical executions.