CAU Mountain&Sea 2026 Humanoid Soccer Small Size Extended Abstract

Bei Liu, Dian Qiu, Yuzu Jiang, Pengyu Jiao, 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 and intelligent multi-agent collaboration. Distinguishing our approach from traditional methods, we introduce two key innovations: (1) **Deep Reinforcement Learning (DRL)** frameworks designed to synthesize dynamic kicking and passing maneuvers that integrate balance control with ball manipulation; and (2) **Graph Neural Network (GNN)** based policies that model the team's spatial topology to enable implicit tactical coordination. These advancements address the challenges of high-dynamic mobility and complex cooperative gameplay in competitive humanoid soccer.

1 Lessons We Learned

1. **Graph-Based Representation for Tactical Awareness**: We learned that fixed role assignment (e.g., rigid Striker/Defender switches) is insufficient for fluid gameplay. Instead, modeling the field as a dynamic graph—where robots are nodes and their interactions are edges—provides a superior representation. By employing **Graph Neural Networks (GNNs)**, our agents can effectively aggregate information from neighbors (teammates) to understand the global tactical context. This allows for emergent cooperative behaviors, such as creating passing triangles, rather than relying on prescripted state machines.
2. **End-to-End Learning for Dynamic Actions**: Previously, we relied on engineered stability criteria (ZMP) for kicking, which limited the robot's ability to kick while moving. We discovered that **Deep Reinforcement Learning (DRL)** allows the robot to learn the complex coupling between support-leg balancing and swing-leg trajectory generation. This shift has taught us that formulating the kick and pass as a unified whole-body control problem—rewarded by ball velocity and directional accuracy—yields significantly more agile and robust actions than separating locomotion and manipulation.

2 Major Problems

1. **Precision and Stability in Dynamic Kicking via DRL**: A critical challenge is enabling the robot to perform accurate passes or powerful shots

without stopping its walking gait. Traditional kinematic planners often fail to handle the sudden momentum shift during a kick. We are currently addressing the "Sim-to-Real" gap in our DRL policies. The core problem is designing a reward function that balances two conflicting objectives: maximizing the *impulse* imparted to the ball for speed/distance, while minimizing the *angular momentum disturbance* to the robot's trunk to prevent falling. Furthermore, the policy must learn to adapt the foot's impact pose to achieve precise passing vectors despite sensor noise.
2. **Scalable Coordination with Graph Neural Networks**: In dynamic environments, the number of relevant agents (teammates in view, opponents blocking paths) varies constantly. Standard fixed-input neural networks struggle with this variability. Our major challenge lies in training GNN-based value functions that can generalize across different team configurations. We aim to solve the credit assignment problem in multi-agent settings, ensuring that the GNN correctly identifies which specific cooperative edge (e.g., a critical pass between two specific nodes) contributed most to a goal, thereby optimizing the team's graph topology for better ball circulation.

3 Plans for RoboCup 2026

1. **GNN-Driven Collaborative Strategy Framework**: We plan to fully deploy a Graph Convolutional Network (GCN) architecture for high-level decision making. In this framework, the state of the game will be represented as a fully connected graph. The GNN will perform message passing to propagate intent (e.g., "I am open for a pass") across the network. This will enable sophisticated collaborative plays, such as *predictive passing*, where the ball carrier passes not to where a teammate *is*, but to where the GNN predicts the teammate *will be*, optimizing space utilization against opponent defense.
2. **Hierarchical DRL for Agile Ball Manipulation**: We will upgrade our motion engine to a hierarchical DRL system. The low-level policy will focus on robust joint actuation and balance recovery, while the high-level policy will learn specific "skills" such as *curve kicks*, *chip passes*, and *intercepting moving balls*. By training these policies in parallel, we aim to achieve a seamless transition between running and kicking. Specifically, we will focus on learning an "action-aware gait" where the robot automatically adjusts its final steps before the kick to maximize foot-to-ball contact accuracy, ensuring high-speed passes and shots.

4 Current Status

We have developed a high-performance humanoid platform optimized for learning-based control. Our current software stack has successfully integrated a **Deep Reinforcement Learning** module for basic omnidirectional walking and static kicking, demonstrating superior disturbance rejection compared to model-based controllers.

On the strategic front, we have established a preliminary **Graph Neural Network** inference engine that runs in real-time on the robot's onboard computer. This engine currently handles simple formation control by treating robot positions as graph nodes, allowing the team to maintain optimal spacing dynamically.

With these foundations in DRL-based motion and GNN-based perception, we are well-prepared for RoboCup 2026, creating a team that not only moves with agile, learned behaviors but also thinks as a unified, graph-connected intelligence.