

SJTU Cyber Shanghai 2026 AdultSize Team

Description

Muye Zhang, Xiaoxiao Zhu, Pengzhi Chu, Chuntao Leng

Shanghai Jiao Tong University, Shanghai, China

zhangmuye@sjtu.edu.cn

Abstract

This paper introduces **SJTU Cyber Shanghai**, a new entrant to the RoboCup 2026 Humanoid League AdultSize competition. Our team originates from the **Cyber Shanghai Club** at Shanghai Jiao Tong University (SJTU)—rated as the university's top-tier "5-Star Academic Club" in 2025. Comprising undergraduate to PhD students focused on embodied AI and humanoid robotics, the team is supported by the China Association for Science and Technology's pilot plan and the SJTU "Qinghe" Innovation Plan. We operate under the **Student Innovation Center**, leveraging its 20-year legacy in RoboCup and the guidance of a former RoboCup World Champion. Building on a track record that includes a Global Top 8 finish (Group 1st) at the inaugural World Humanoid Robot Games and a National Championship in the Humanoid Robot Relay, we aim to transfer our agility to the soccer gameplay. We utilize the **MagicBot Z1** as our hardware platform, featuring a custom-engineered head module designed to optimize the field of view for match conditions. Our proprietary software stack integrates **YOLOv8** for robust visual perception with a comprehensive **Sim-to-Real** control pipeline. By combining **Reinforcement Learning (RL)** and **Imitation Learning**, we have developed stable models for omnidirectional running, dynamic kicking, and autonomous fall recovery. Furthermore, we implement a self-correcting localization model that fuses vision and IMU data to mitigate drift, ensuring high-performance autonomous gameplay.

1 Introduction

SJTU Cyber Shanghai is an interdisciplinary team selected from the SJTU Cyber Shanghai academic club, **the highest-rated (5-star) academic club** at Shanghai Jiao Tong University. Our members span undergraduate, master's, and doctoral levels, specializing in humanoid robotics, embodied intelligence, and digital media. The team is incubated by the **Leng Chuntao Robot Innovation Studio** and the

Embodied AI Makerspace within the SJTU Student Innovation Center. Our work is supported by the "Qinghe" Innovation Plan and the China Association for Science and Technology's initial pilot program.



While 2026 marks our debut in the AdultSize League, our team stands on a foundation of extensive competitive success and public recognition.

- **Competitive Heritage:** As the sole university representative from Shanghai at the first World Humanoid Robot Games, we achieved a Global Top 8 finish. Domestically, we have secured the National Championship in the Humanoid Robot Relay and multiple First Prizes at the China Robot Competition for two consecutive years.
- **RoboCup Legacy:** We benefit from the SJTU Student Innovation Center's 20-year history of RoboCup participation. Our technical roadmap is directly guided by **Chuntao Leng** (Chief Engineer) and **Xiaoxiao Zhu**, who possess a distinguished track record in the competition. They previously led SJTU teams to achieve the **2017 RoboCup Small Size League World Championship**, as well as the **China National Championships** in both the **Standard Platform League** and **Middle Size League**. Leveraging this world-champion caliber experience, we utilize their expertise to accelerate our entry into the AdultSize league.
- **Social Impact:** Our fusion of technology and culture has garnered significant attention from **authoritative national media outlets**, including **CCTV, China News**

Service, The Paper, The Beijing News, and Dragon TV. Beyond media recognition, we have organized over 100 science outreach activities, serving tens of thousands of individuals.



Technically, our approach bridges the gap between frontier Embodied AI research and competitive reliability. Leveraging the **MagicBot Z1** as our kinematic foundation, we have engineered a proprietary **high-fidelity visual perception system** optimized for the dynamic occlusion and lighting conditions of soccer environments.

Our control architecture represents a paradigm shift from classical model-based methods to a **data-driven Deep Reinforcement Learning (DRL) framework**. By integrating **Imitation Learning (IL)** with motion capture retargeting, we train whole-body control policies across a suite of advanced simulation environments, including **Isaac Gym, Isaac Lab, and MuJoCo**. This multi-platform approach ensures our policies—capable of robust omnidirectional locomotion, dynamic kicking, and **autonomous fall recovery**—are validated against diverse physics engines before deployment. To further guarantee seamless **Sim-to-Real transfer**, we employ extensive **Domain Randomization** to adapt to physical uncertainties.

Simultaneously, our perception stack utilizes a fine-tuned **YOLOv8** backbone for real-time semantic feature extraction. To address the challenge of cumulative drift in dead-reckoning, we deployed a **learning-based state estimation module** that fuses

categories of field markers.

First are the corner markers. We utilize a fine-tuned **YOLOV8** model for object detection. As shown in the figure below, our system successfully detects field features such as the "L-Cross" 、 "T-Cross" and "Ball" with high confidence and estimates their depth in millimeters.

Meanwhile, we also identify the white field lines. First, we use a scanline method to locate potential points lying on white lines, then further filter these to retain more reliable line points, which are subsequently connected into straight segments or arcs. This category of markers provides valuable supplementary information, especially in situations where the first category of points is not clearly visible.



- **Localization:** We have implemented a coordinate transformation system that maps camera-space detections to the robot's base footprint and the global map.

Based on the provided images and detailed technical description, here is the refined and expanded **Vision and Localization** section for your TDP. This version articulates the specific algorithms (MCPF), landmark types, and the adaptive fusion strategy, supported by the quantitative results from your images.

3.1.1 Adaptive Monte Carlo Localization (AMCL)

The core localization engine is a particle filter that fuses sensor data through a three-step process:

1. **Prediction:** Particles are propagated using the motion model derived from the odometer and IMU.

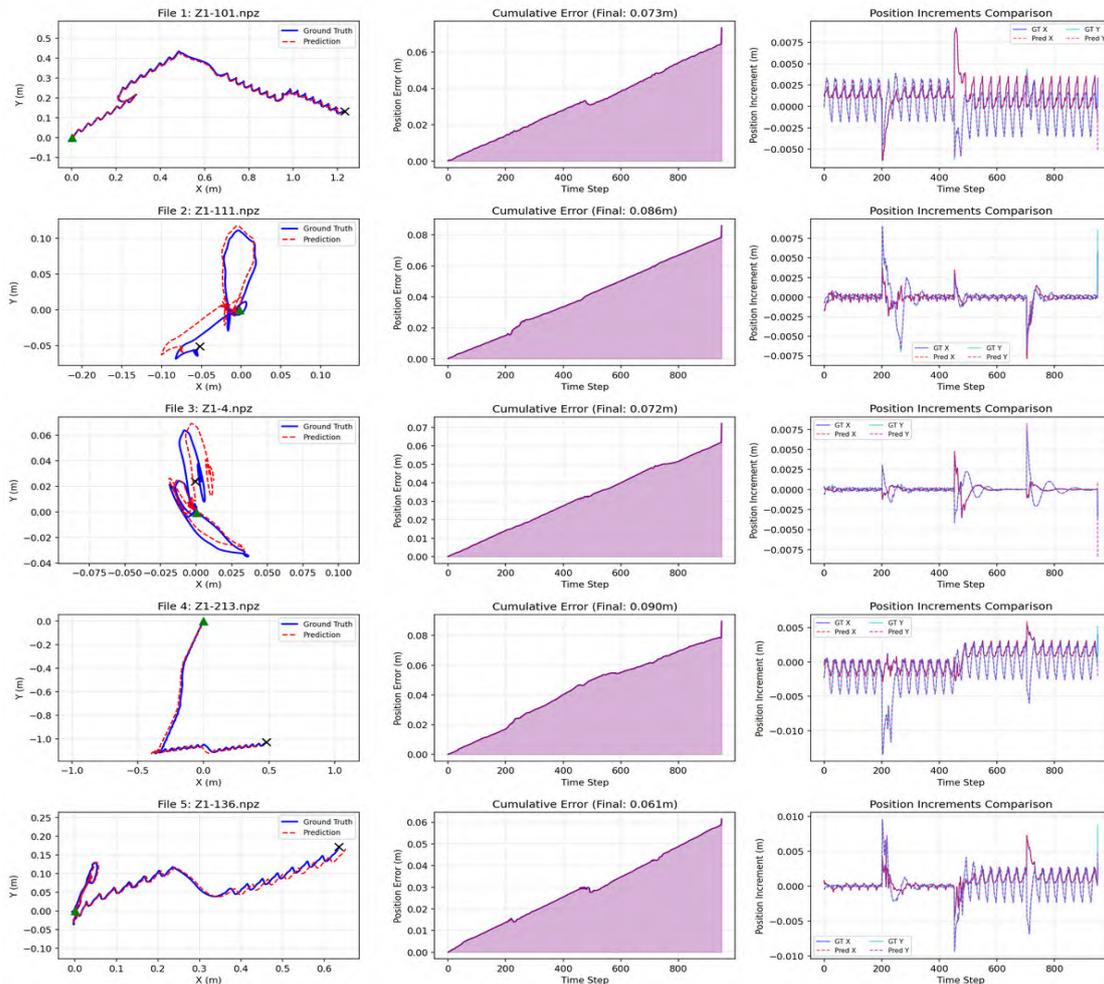
2. **Update:** Weights are updated based on the residuals between the predicted and observed positions of the visual landmarks (X, P, T, L types, and Line segments).
3. **Resampling:** The system focuses on high-probability particles to converge on the robot's true pose.

To ensure robustness, we employ an **Adaptive Fusion Strategy**. The system dynamically adjusts sensor weights based on environmental conditions. In scenarios with visual occlusion, the weight of the odometry is increased to maintain continuity. Conversely, when visual features are abundant, the system prioritizes visual data to eliminate drift. A specialized **Fall Recovery Mode** utilizes IMU gravity vectors to rapidly re-initialize the particle cloud, ensuring the robot can relocalize instantly after standing up.

3.1.2 Trajectory Optimization Results

To validate the system, we tested the "AutoOdom" module, which corrects raw odometry using the learned error models. As shown in figure below, the predicted trajectory (Red Dashed Line) closely tracks the Ground Truth (Blue Solid Line) across complex movement patterns. Quantitative analysis confirms the system's precision, achieving an **Average RMSE of 0.000210** and an **Average Final Position Error of only 0.0765m** over long-duration tests.

AutoOdom Stage 1 Visualization Results
 Avg RMSE: 0.000210 | Avg Final Error: 0.0765m | Avg Relative Error: 6.83%

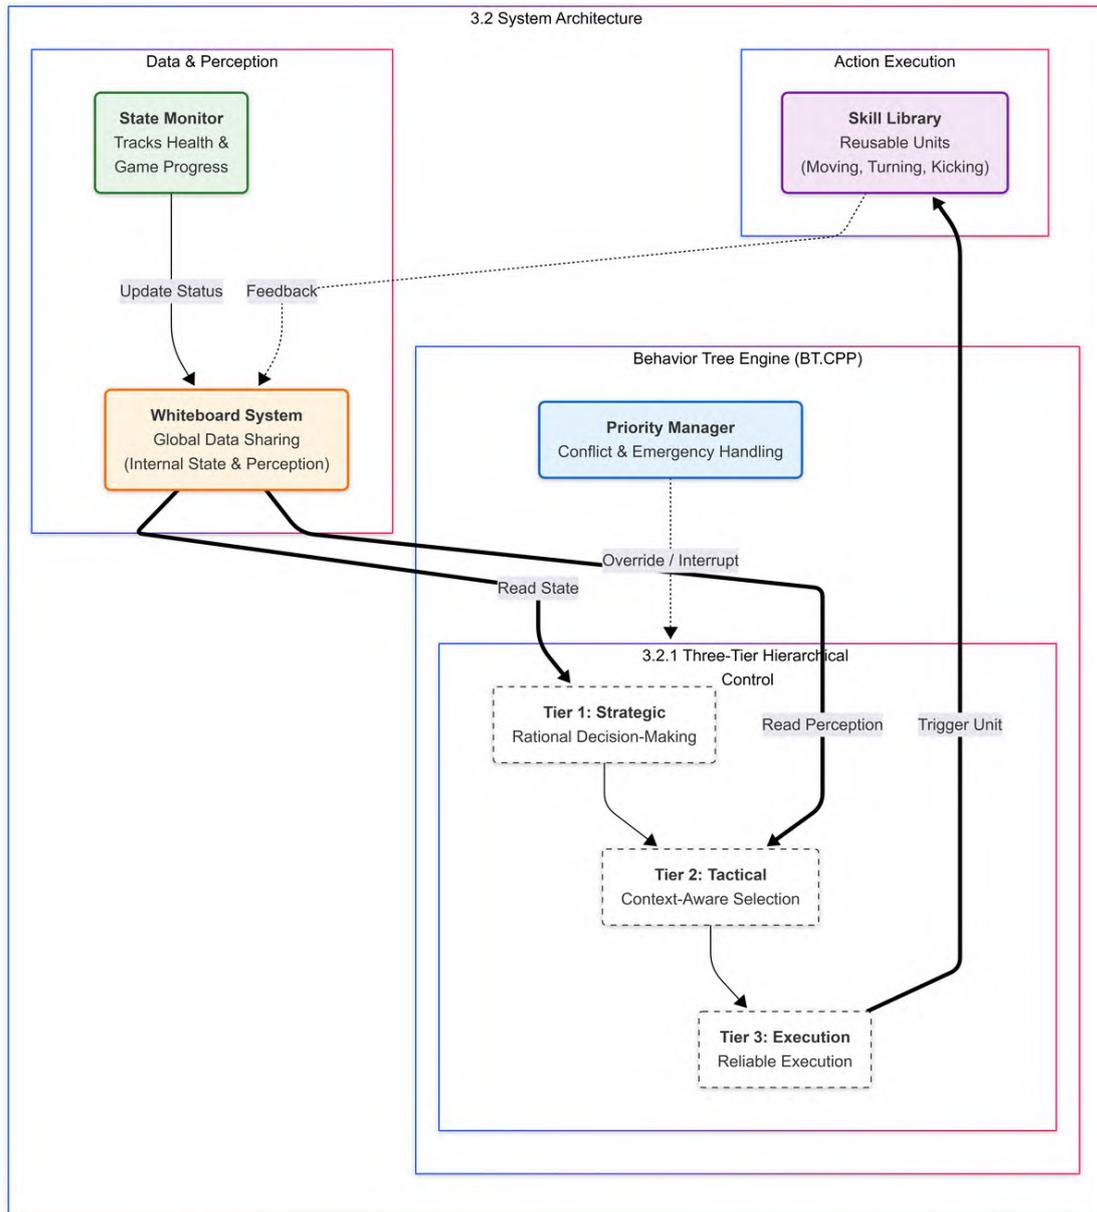


- **Odometry Correction:** To mitigate drift, we trained a model to autonomously correct IMU and odometry data. Preliminary results demonstrate that our prediction model (red dashed line) closely follows the ground truth (blue line), significantly reducing cumulative position error compared to raw sensor data.

3.2 Behavior Control System

Our robot employs a modular behavior control system based on Behavior Trees (BT). This system utilizes a hierarchical, reactive decision-making framework to achieve comprehensive behavior management—spanning from low-level motion control to high-level game strategy. By integrating the robot's real-time perceptual information, match status data, and predefined tactical strategies, the system generates optimal behavior

sequences tailored to the current situational context.



3.2.1 Hierarchical Control Structure

The behavior control system adopts a three-tier hierarchical architecture to ensure rational decision-making and reliable execution:



3.2.2 Core Components

Behavior Tree Engine: Implements the BT.CPP framework based on the ros-humble-behaviortree-cpp package, providing efficient behavior tree parsing and execution capabilities.

Whiteboard System: A global data-sharing space that stores the robot's internal state and external perceptions.

Skill Library: A repository of reusable basic behavior units, such as moving, turning, kicking, etc.

State Monitor: Tracks the robot's health status and game progress in real-time.

Priority Manager: Handles behavior conflicts and emergency response.

3.3 Locomotion via Imitation and Reinforcement Learning

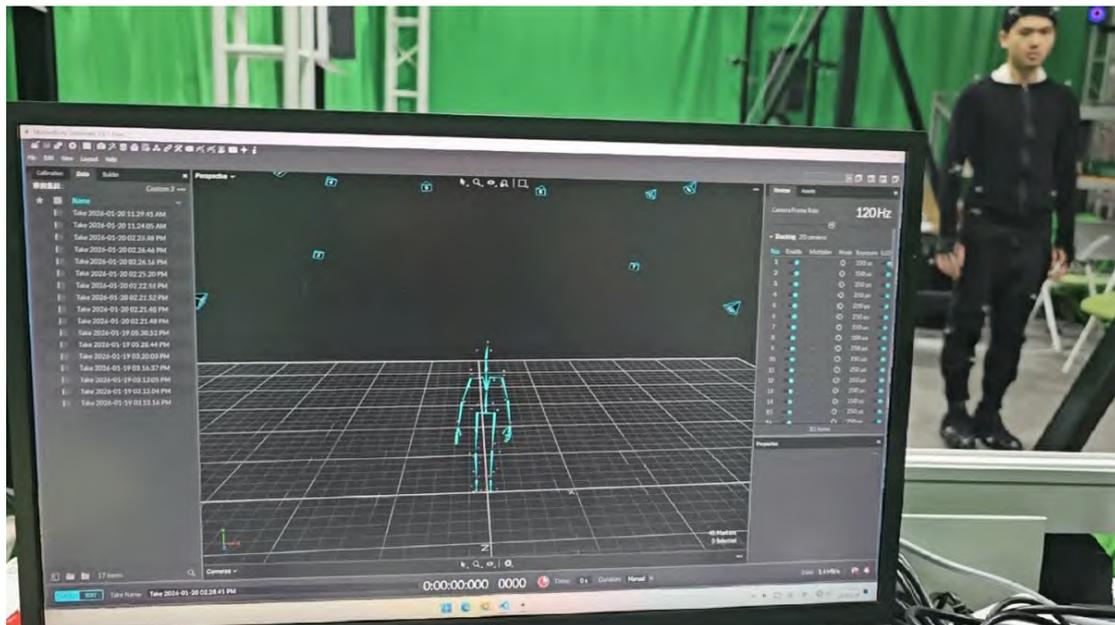
Our locomotion architecture represents a paradigm shift from classical model-based control to a data-driven **Reinforcement Learning (RL)** framework. This approach enables the MagicBot Z1 to exhibit dynamic, robust, and human-like agility that is difficult to hard-code manually.

3.3.1 Motion Capture and Retargeting Pipeline

To bootstrap our learning process with bio-plausible motions, we

established a proprietary motion capture workflow rather than relying solely on open-source datasets.

- **Data Acquisition:** We utilized an optical motion capture system to record a custom dataset of human soccer players performing task-specific maneuvers, including omnidirectional running, sudden stops, side-kicks, and defensive saves.



- **Kinematic Retargeting:** Since the kinematic chain of a human differs from the MagicBot Z1, we apply an offline optimization process. This retargeting minimizes the geometric difference between key end-effectors (feet, hands, head) while respecting the robot's joint limits and URDF constraints, generating a reference motion library for the RL agent to imitate

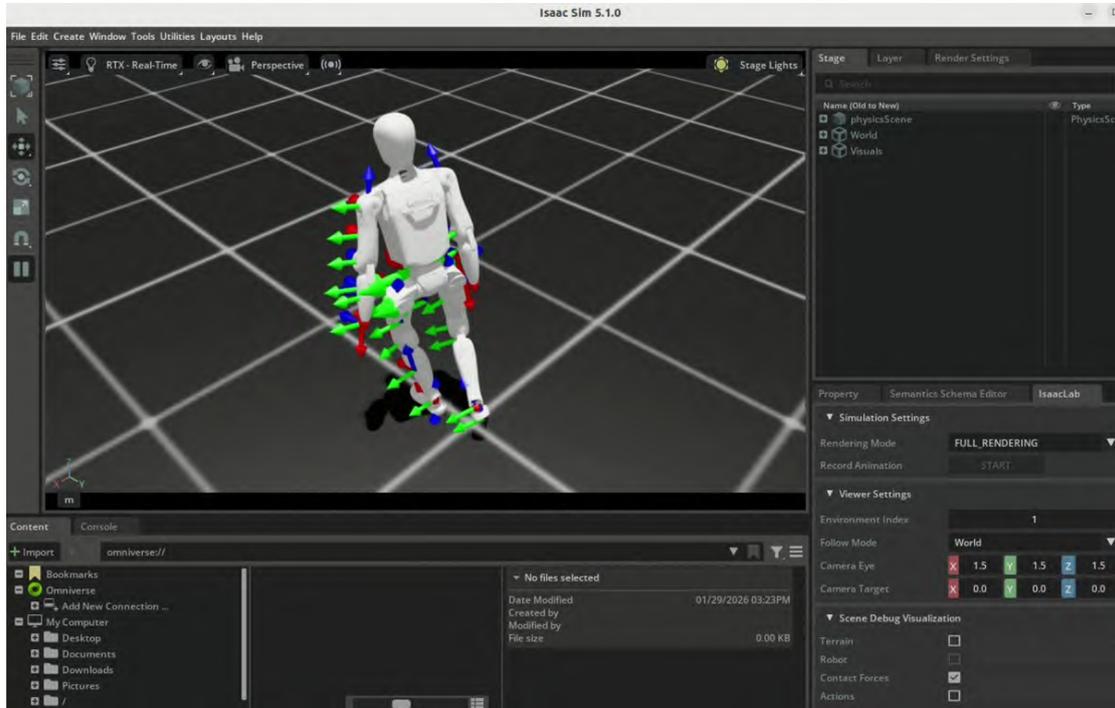
3.3.2 Policy Training and Domain Randomization

We train our policies using PPO (Proximal Policy Optimization) within massively parallel simulators (Isaac Gym and Isaac Lab). To bridge the "Sim-to-Real" gap, we employ extensive Domain Randomization:

- **Physical Parameters:** We randomize link masses, joint friction, and damping coefficients during training to prevent the policy from overfitting to the perfect simulation physics.
- **Observation Noise:** Gaussian noise is added to the simulated IMU and

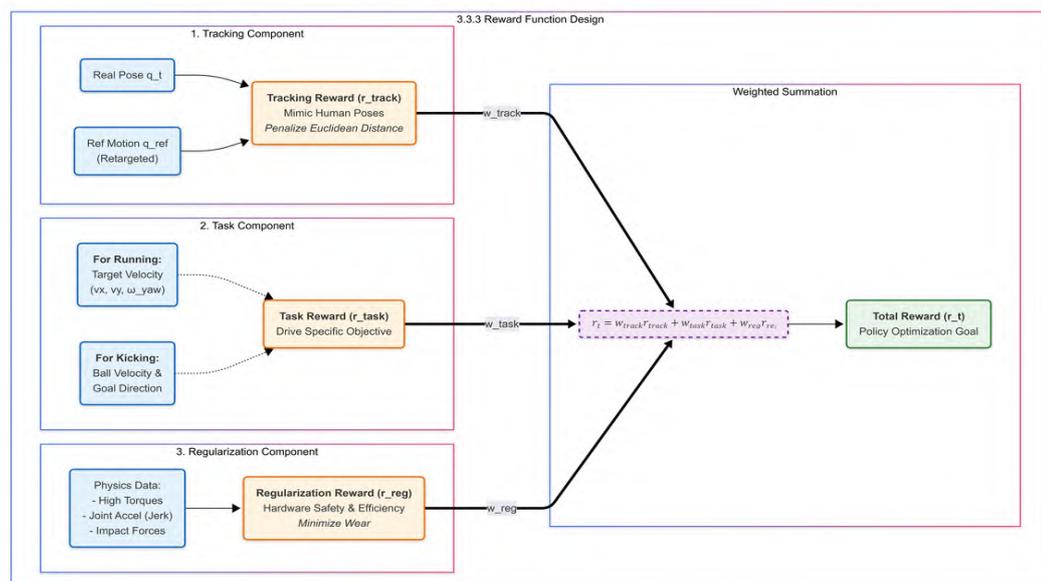
joint encoder readings to mimic real-world sensor imperfections.

- **Perturbations:** We apply random external force vectors to the robot's torso during training, ensuring the policy learns active push recovery.



3.3.3 Reward Function Design

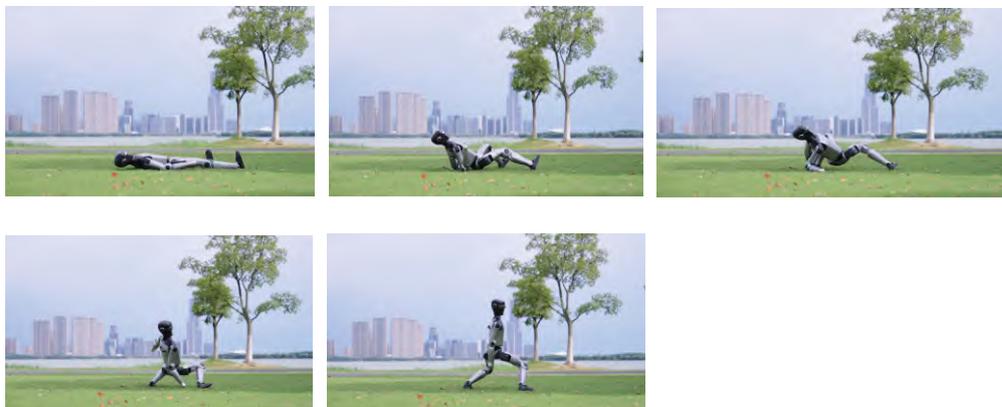
The reward function is designed as a weighted sum of three primary components: **Tracking**, **Task**, and **Regularization**.



- **Tracking Reward:** Encourages the agent to mimic the retargeted human poses. It penalizes the Euclidean distance between the robot's joint positions and the reference motion .
- **Task Reward:** Drives the specific objective.
 - *For Running:* It rewards tracking the target velocity command.
 - *For Kicking:* It rewards the velocity of the ball relative to the goal direction upon contact.
- **Regularization Reward:** Critical for hardware safety. We penalize high joint torques, excessive joint accelerations (jerk), and foot impact forces to ensure smooth, energy-efficient motion that minimizes wear on the actuators.

3.3.4 Skill Set and Status

- **Omnidirectional Running:** The agent has mastered a stable gait capable of tracking velocity commands while rejecting external disturbances.
- **Dynamic Kicking:** Unlike static kicking engines, our RL policy enables **dynamic kicking** (kicking while moving) and specialized **side-kicks**, allowing for faster gameplay transitions.
- **Fall Recovery:** We developed a specialized T1FallRecovery task. The policy takes the robot's current orientation (face-up or face-down) as input and outputs a sequence of actions to return to a standing posture, validated to work from arbitrary initial states.

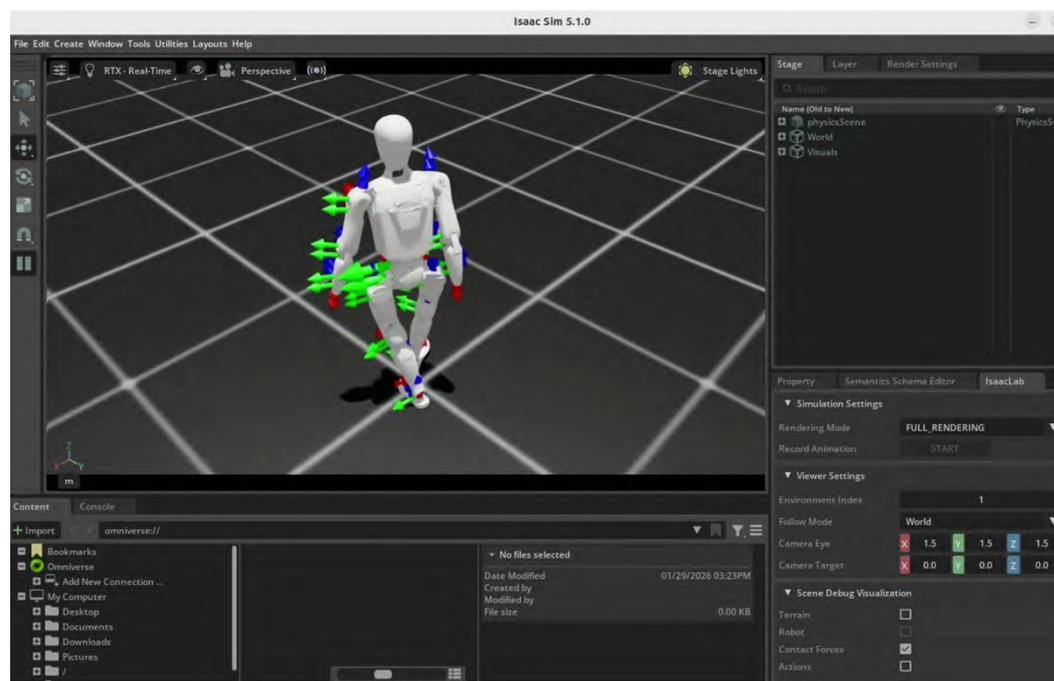


- **Goalkeeping:** We are training a position-conditioned policy where the

robot observes the ball's trajectory and executes defensive blocks (squatting or jumping) to intercept the ball, based on the "Humanoid Goalkeeper" research framework.

- **Advanced Football Techniques:** We leverage a hybrid of **Imitation Learning** and **Reinforcement Learning** to master high-difficulty soccer-specific maneuvers. This portfolio includes:

- **Tackling:** Sliding maneuvers to intercept the ball from opponents effectively.
- **Long-Range Clearance:** Optimized power kicking motions ("Big Foot") designed to clear the ball from defensive zones to the opposite field.
- **Lob Shots:** High-trajectory kicks (High Balls) enabling the robot to lift the ball over obstacles or defenders



3.4 Strategy and Real-Time Tactical Interface

To effectively orchestrate the RL-based locomotion skills described above, we have developed a high-level strategic framework supported by a custom visualization tool. This system bridges the gap between low-level control and high-level team play.



3.4.1 Proprietary Tactical Board

We developed a real-time, web-based tactical board (currently for internal development) to facilitate strategy design and monitoring. This interface serves two critical functions:

- **Real-Time Visualization:** The system ingests high-precision pose estimates from our Vision-Odometry Fusion system (Section 3.1). It renders the live positions of teammates (depicted as Blue nodes, e.g., Player 1, 2, 3) and perceived opponents (Red nodes, e.g., Player 4, 5, 6) on a virtual 2D field.
- **Human-to-Robot Tactic Transfer:** The interface is designed to transfer human soccer coaching knowledge to autonomous agents. It allows us to simulate and refine standard tactical plays—such as formation adjustments or defensive alignment—ensuring they are geometrically feasible for the robots.

3.4.2 Strategic Execution

The insights derived from the tactical board are translated into a hierarchical behavior tree that governs the robot's autonomy.

- **Context-Aware Action Selection:** Based on the global game state visualized in the interface, the system dynamically triggers the appropriate

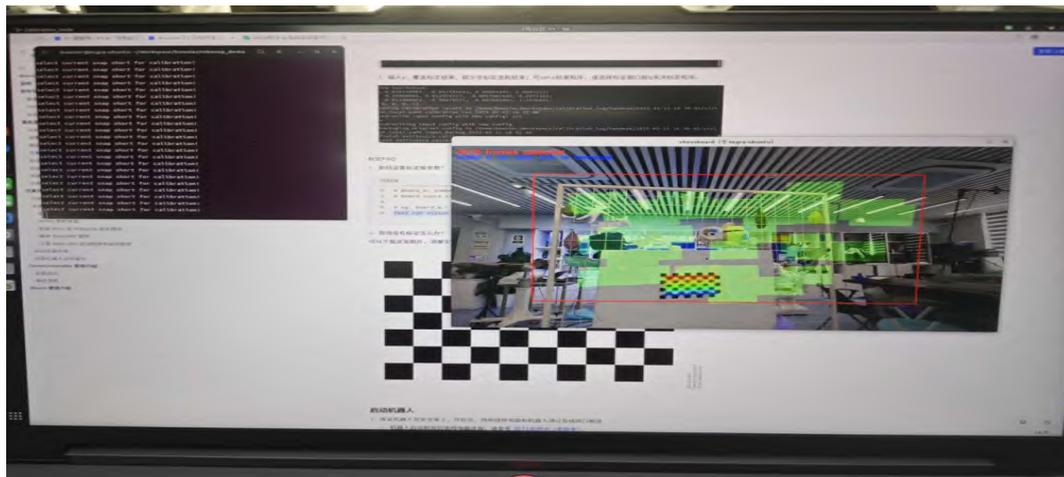
RL policy (Section 3.2).

- **Tactic Triggers:** As shown in the interface control panel, specific high-level directives such as "Run Tactic," "Shoot Tactic," or "Clear Tactic" can be tested and tuned. For instance, a "Shoot" command will activate the navigation stack to align the robot with the goal using the *Omnidirectional Running* policy, followed by the *Dynamic Kicking* policy upon ball contact.

4 Current Status and Experiments

We are currently in the final stages of Sim-to-Real deployment.

- **Vision:** The calibration algorithms for the camera and joints are operational, and the TF tree correctly links the camera_optical_frame to the base_footprint.



- **Control:** The RL running and fall recovery policies have been successfully validated in simulation and are being deployed to hardware.



5 Conclusion

In this paper, we introduced **SJTU Cyber Shanghai**, a new team aiming to bring fresh energy and innovation to the RoboCup 2026 AdultSize League. Our entry represents a significant step in translating our academic background into robust functionality. By rigorously testing a **Sim-to-Real RL pipeline** and a custom **visual perception module** on the **MagicBot Z1**, we have laid a solid foundation for autonomous gameplay.

Our initial results in omnidirectional locomotion and fall recovery demonstrate the **potential** of our data-driven approach to handle the dynamic demands of soccer. However, we recognize that competitive robustness requires continuous iteration. We view this debut participation as a valuable opportunity to test the limits of our algorithms against real-world uncertainties.

We are honored to join this community and are committed to sharing our development progress. We look forward to competing alongside and learning from veteran teams, striving to grow into a strong contributor to the Humanoid League.

References

- [1] Yushi Wang, C. Luo, P. Chen, et al. "Learning Vision-Driven Reactive Soccer Skills for Humanoid Robots". In: arXiv:2511.03996 [cs.RO] (November 2025)