BahiaRT 2026: Team Description Paper for RoboCup Humanoid Soccer League

Alan Vitor Nascimento, Luis Henrique Cunha, Caio Lucena, Melissa Damasceno, Pedro Rabelo, Bruno Simões, Felipe Silva, Janderson Blanski, João Cleber Pinto, Ma Qi Chao, Mirella Costa, Rafael Mattos, Sara Adriana Oliveira, Ana Patricia Magalhães, Jorge Campos, José Grimaldo Filho, Josemar Rodrigues de Souza, and Marco A C Simões

Bahia State University (UNEB/ACSO), Salvador, BA, Brazil
teambahiart@gmail.com
http://www.acso.uneb.br/bahiart/

Abstract. BahiaRT proposes joining the RoboCup Humanoid Soccer League in 2026. With 20 years of experience in the 3D Soccer Simulation League (11 regional championships, consistent World Cup participation), we are transitioning to physical humanoid robotics with a research focus on systematic sim-to-real transfer. We have acquired two Booster T1 robots and demonstrated initial capabilities at October 2025 events. Our modular architecture enables transferring reinforcement learning skills and multi-agent strategies from simulation to physical robots. This paper describes our technical approach, current progress, and research goals for HSL competition.

1 Introduction

The RoboCup Humanoid Soccer League requires integration of perception, locomotion, multi-agent coordination, and real-time decision-making under dynamic conditions. Transferring capabilities from simulation to physical robots remains challenging due to the reality gap—discrepancies in dynamics, sensing, and actuation between simulated and real systems.

BahiaRT has participated in the 3D Soccer Simulation League for nearly two decades, achieving 11 regional championships and consistent World Cup qualification since 2009. In 2025, we acquired two Booster T1 robots and conducted initial demonstrations, marking our transition to physical humanoid robotics. This paper presents our systematic sim-to-real approach, combining modular architecture (Section 3), transfer methodologies (Section 4), and multi-agent strategy adaptation (Section 5), alongside our community contributions (Section 2).

2   Team History and Contributions to the RoboCup Community

BahiaRT has been participating in RoboCup competitions for nearly 20 years since 2007, primarily in the 3D Soccer Simulation League where the team achieved 11 regional championships in Brazil and Latin America, and consistent World Cup qualification since 2009. Most recently, the team was crowned Champion of the Flying Robots Demo in 2025. In October 2025, BahiaRT conducted its first physical humanoid robot demonstrations at the RoboCup Brazil Open and AI Brazil Experience using Booster T1 robots, marking the transition from simulation to real-world robotics.

The team's research contributions focus on reinforcement learning and multi-agent coordination. Notable publications include set-play selection for coordinated moves [2], Learning from Demonstration (LfD) [7, 5], multi-agent passing and defense strategies [4], and most recently at RoboCup Symposium 2025, a paper on reinforcement learning for robust kicking skills in the NAO robot. The team has also contributed to service robotics research [1, 3].

BahiaRT has developed the MR-SoccerServer [8], BahiaRT Gym [6], and is actively collaborating on the new MuJoCo-based Simulator for the 3D League, providing a realistic platform for testing strategies transferable to physical humanoid soccer.

Team members have served on the Soccer Simulation League Executive Committee (2015-2022), organized RoboCup 2025 in Salvador, and held RoboCup Brazil leadership positions.

3   BahiaRT 2026 Architecture

Leveraging our extensive experience in the 3D Soccer Simulation League, we designed a modular system architecture that enables systematic sim-to-real transfer of both low-level skills and high-level strategies. This design provides generic and decoupled packages for code reuse across simulation and physical robots, as illustrated in Figure 1.
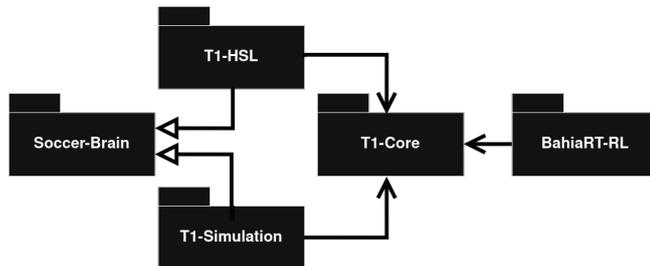

Fig. 1: High-level architecture of the BahiaRT modules.

Soccer-Brain serves as the core decision-making module, implementing high-level multi-agent reasoning including passing coordination, defensive marking, set plays, and team formations. It operates on an abstract world model with low-level control decoupled via interfaces, ensuring strategy reuse across different execution environments and robot platforms.

T1-Core provides robot-specific components for the Booster T1, including URDF description files, forward/inverse kinematics, and controller implementations. It hosts trained neural network policies and exposes standardized observation and action interfaces, enabling the same learned skills to be deployed in both simulation and physical robots.

T1-HSL integrates Soccer-Brain and T1-Core into a distributed ROS2-based ecosystem for real-robot deployment. This layer implements environment-specific components: vision algorithms for ball and robot detection, multi-agent communication via WiFi, interfaces with the official game controller, and integration with the T1 SDK and hardware drivers. It orchestrates the complete system execution on physical robots during matches.

T1-Simulation implements the same Soccer-Brain logic within a serial execution model for the 3D Soccer Simulation League. It communicates with the simulator via socket-based interfaces while relying on T1-Core for skill execution, maintaining behavioral consistency with the physical system.

BahiaRT-RL serves as our primary skill acquisition framework, providing robot-, simulator-, and RL-framework-agnostic interfaces for scalable training on both CPU and GPU clusters. This design enables systematic development and validation of skills in simulation before transfer to physical robots, which is central to our sim-to-real research approach.

4 Sim-to-Real Transfer

We have acquired two Booster T1 robots and conducted initial deployment experiments at demonstration events in October 2025 (RoboCup Brazil Open and AI Brazil Experience). Current capabilities include: stable omnidirectional walk at speeds comparable to the league state-of-the-art, short-range kick (2-3 meters) integrated with directional approach, efficient get-up recovery from both prone and supine positions, and vision-based detection of ball, goal posts, field lines, teammates, opponents, and referees. Self-localization on the field is also operational. Video demonstrations of these skills will be submitted with this TDP.

Our sim-to-real methodology combines three complementary strategies to address the reality gap. **Domain Randomization** is applied during reinforcement learning training to motor noise, ground friction, actuation delays, and sensor latency, promoting policy robustness to physical uncertainties. **System Identification** calibrates simulator parameters—including joint limits, PID controller gains, center of mass offsets, and contact dynamics—using data collected from real robot teleoperation and motion analysis. **Model-based control priors** derived from inverse dynamics (Euler-Lagrange formulation) constrain the RL

search space and ensure physically plausible motions, reducing the sim-to-real gap for dynamic behaviors.

This approach enables transferring simulation skills—walk, sprint, dribbling—while refining them with real-world feedback.

5   High-Level Strategy Transfer: Coordination and Setplays

Setplays are our main competitive differentiator, developed and refined over many years in the 3D Soccer Simulation League, with strong prospects for transferring this expertise to the Humanoid Soccer League. Our team has extensive experience developing coordinated strategies for both dead-ball situations (offensive and defensive kickoff formations, corner kicks, free kicks, and goal kicks) and open-play scenarios (marking schemes, passing sequences, and defensive positioning).

Our approach leverages multi-agent communication and dynamic role allocation, enabling robots to execute collective plays rather than isolated actions. The effectiveness of setplays as high-level team behaviors that encode human tactical knowledge has been formally investigated within our team [9, 2, 4], demonstrating gains in collective performance and tactical consistency.

For the Humanoid Soccer League, we are adapting these strategies by incorporating contact-aware positioning and balance-preserving motion primitives, addressing the physical constraints of real robots while preserving the tactical sophistication developed in simulation.

6   Conclusion and Future Work

This paper presented BahiaRT's proposal for the Humanoid Soccer League in 2026. After 20 years in the 3D Soccer Simulation League, we are transitioning to physical robotics with a focus on systematic sim-to-real transfer. Our modular architecture enables code reuse between platforms, while our methodology, combining domain randomization, system identification, and model-based control, addresses the reality gap. We have demonstrated initial capabilities on two Booster T1 robots, including stable locomotion, vision-based perception, and self-localization.

Our immediate goals for RoboCup 2026 include developing a long-range kick through reinforcement learning, transferring positioning and passing strategies from simulation, and integrating vision with learned locomotion policies. We aim to contribute to the league by validating sim-to-real methodologies and sharing open-source tools that can accelerate research across all teams.

References

1. Rodowanski, I.J., Mascarenhas, A.P.F.M., Simões, M.A.C., Campos, J.A.P.d., Silva Filho, J.G.d., Souza, J.R.d.: A new architecture to bill service robot. In:

Proceedings of the V Brazilian Humanoid Robot Workshop (BRAHUR) and VI Brazilian Workshop on Service Robotics (BRASERO). Curitiba (2024)

2. Santos, C.S., Mascarenhas, A.P.F.M., de Souza, J.R., Simões, M.A.C.: Load and selection policies for set-plays in the robocup 3d soccer simulation league. In: 2025 Brazilian Conference on Robotics (CROS). pp. 1–6. Belo Horizonte (2025)

3. da Silva, R.M., da Silva, L.Q., de Souza, J.R., Simões, M.A.C., Campos, J.A.P., Mascarenhas, A.P.F.M.: Petri net-based modelling for autonomous service robots. In: III Brazilian Workshop on Service Robotics (BRASERO). pp. 41–46. São Bernardo do Campo (2019)

4. Simoes, M.A.C., et al.: Strategy planner: Enhancements to support better defense and pass strategies within an lfd approach. In: 2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC). pp. 46–52. IEEE (2020)

5. Simões, M.A.C., et al.: Generating a dataset for learning setplays from demonstration. SN Applied Sciences 3, 608 (2021)

6. Simões, M.A.C., et al.: Bahiart setplays collecting toolkit and bahiart gym. Software Impacts 14, 100401 (2022)

7. Simões, M.A.C., da Silva, R.M., Nogueira, T.: A dataset schema for cooperative learning from demonstration in multi-robot systems. Journal of Intelligent & Robotic Systems pp. 1–20 (2019)

8. Simões, M.A.C., Souza, J.R., Pimentel, F.A.M., Frias, D.: Mr-simulator: A simulator for the mixed reality competition of robocup. In: del Solar, J.R., Chown, E., Plöger, P.G. (eds.) RoboCup 2010: Robot Soccer World Cup XIV, vol. 6556, pp. 82–96. Springer, Berlin / Heidelberg (2011)

9. Simões, M.A.C.: Learning from demonstration of coordinated plans in multi-agent systems. Ph.D. thesis, Institute of Computing, Federal University of Bahia (2022)