Towards the ultimate goal of RoboCup with a standard humanoid robot platform

Mingguo Zhao
Departmant of Automation
Tsinghua University
Beijing, China
mgzhao@tsinghua.edu.cn

Hao Dong
Booster Robotics
Beijing, China
donghao@boosterobotics.com

Xinyu Han
Booster Robotics
Beijing, China
hanxinyu@boosterobotics.com

Qi Li
Booster Robotics
Beijing, China
liqi@boosterobotics.com

***Abstract*—We discuss the significance of utilizing a standard humanoid robot platform to achieve the ultimate goals of the RoboCup initiative. Additionally, we provide a detailed overview of our Teen-Sized humanoid robot, Booster T1, along with a description of our RL-based locomotion framework. Finally, we outline our exciting and ambitious future research plans for humanoid robot soccer to attract more teams to participate and be interested in our work.**

Keywords—RoboCup, Humanoid Robot, Standard Platform

I. INTRODUCTION

The ultimate goal of the RoboCup initiative is to achieve, by the middle of the 21st century, a team of fully autonomous humanoid robot soccer players that can win a soccer match under official FIFA rules against the World Cup champions[1]. According to the roadmap for the humanoid league[2], the 11v11 matches for humanoid robots will commence in a few years. Key research themes for the league include high-speed locomotion, powerful kicks, practical throw-ins, enhanced perceptual abilities, and intelligent teamwork. However, assembling a team of 11 appropriately sized humanoid soccer robots may exceed the budgets of many universities, and training on an entire field poses logistical difficulties due to space constraints. To address these issues, we propose using a standard humanoid robot platform to organize soccer competitions, facilitating cooperation and the development of joint teams in the future. Essential requirements for these standard humanoid robots include advanced perception, mobility, ball control skills, and intelligence.

This paper introduces our innovative Teen-Sized humanoid robot, Booster T1. This robot, designed to embody the intelligence necessary for a soccer robot, solves the challenges of assembling a team of 11 appropriately sized humanoid soccer robots. We will provide a comprehensive understanding of its unique hardware features, mobility, and dexterity, showcasing its potential contributions to robotics and artificial intelligence.

II. DESCRIPTION OF BOOSTER T1

A. *Robot Specifications*

The Booster T1 robot consists of a head, torso, arms, and legs, with 23 degrees of freedom, allowing for flexible movement, as shown in Fig.1. The specifications of the robot are depicted in Table I.

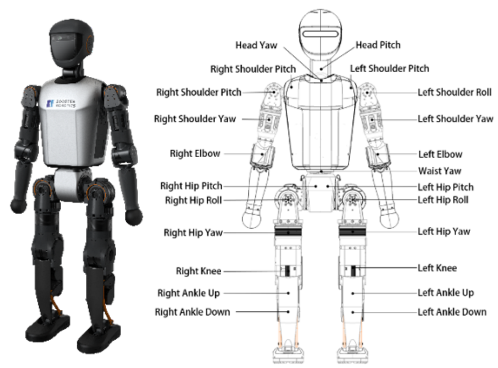

Fig. 1. Booster T1

TABLE I Hardware Specifications

Height(m)	1.18
Weight(kg)	30
DOF	23
Forward Speed(m/s)	1.2
Rotation Speed(rad/s)	1.5
working hour(h)	1.5
charging time(h)	⩽2
CPU	i7 *1
GPU	AGXOrin*1
Camera	binocular depth camera
Audio	Microphone, speakers
Wireless E-Stop	Yes
Wireless network	WIFI 6

The comparison between T1 and NAO， which is currently used in RoboCup SPL, is shown in Table II.

TABLE II Comparison between T1 and NAO

	Booster T1	SoftBank NAO
Height(m)	1.18	0.58
Weight(kg)	30	5.48
DOF	23	25
Forward Speed(m/s)	1.2	0.3
working hour(h)	1.5	1.0
CPU	i7 *1，4.8GHz	ATOM*1，1.9G
GPU	AGXOrin*1,200TOPS	None
Camera	binocular depth camera	2D cameras*2
Joint Control Mode	Position Mode Torque Mode Position Velocity Torque PID Mode	Position Mode

As shown in Table II, Booster T1 has the advantages of computing power, mobility, and joint control, which are the key features for soccer competition with AI.

B. AI-oriented computing system

In recent years, deep learning and LLM approaches have revolutionized many areas of Artificial Intelligence, which requires great computing power and multimodal data processing capability. T1's computing system, OneBox, is specifically designed for algorithm deployment of embodied intelligence. OneBox has three features as follows:

- A variable topology network based on a PCle bus supports the full interconnection of heterogeneous computing units, which includes X86(I7) and GPU(AGX Orin).
- A synchronized mechanism between physical simulation engines and heterogeneous computing platforms using a distributed clock to evaluate hardware-aware algorithms online accelerates the Sim-To-Real deployment process.
- A distributed network based on a high-speed ethernet bus to connect diverse actuators and sensors.

III. Reinforcement Learning Pipeline

We developed a reinforcement learning (RL) pipeline for the Booster T1 robot. This pipeline features a flexible RL training toolkit based on PyTorch and Isaac Gym and a C++ deployment of the policy using Eigen. This enables us to train controllers for locomotion skills and achieve zero-shot sim-to-real transfer.

A. RL Training Toolkit

Our toolkit provides essential utilities for reinforcement learning (RL) algorithms, such as data collection, statistics recording, and classic algorithm functions, without pre-defined training procedures. Users must design their model structures and write training code, but the toolkit is compact and can be mastered in under an hour, enabling the efficient creation of complex training processes. The toolkit includes a simulated training environment based on Isaac Gym, employing domain randomization to address the sim-to-real gap. Users can customize observations, actions, and rewards, making it an ideal starting point for developing RL environments.

B. On-board Deployment

We deploy our policy on a 4.8 GHz Intel i7 CPU, utilizing the Eigen library for matrix computations. This configuration is efficient enough to support real-time inference for small policy models.

C. Push-Recovery Policy

We deployed a push-recovery policy for the Humanoid League's Technical Challenge, effectively withstanding impacts of 22.5 Ns from both front and back, achieving the highest contest score. The training follows the ARMA framework, incorporating a traditional actor-critic model and two encoders for privileged information and observation history. The training consists of three phases:

- Joint training of the actor and privileged information encoder using RL, facilitating quick convergence.
- Fixing the trained actor and privileged information encoder while training the observation history encoder through supervised learning.
- The observation history encoder's output is used as input for the actor, jointly trained with RL before real robot deployment.

The policy encourages adherence to a periodic reference stepping trajectory, allowing dynamic adjustments in stepping speed. Disturbances simulate the impact of a heavy object by altering the robot's linear and angular velocity. Observations include proprioceptive sensor readings and reference trajectory phases, with typical gait periods around 0.7 seconds, decreasing to 0.5 seconds during severe disturbances. The robot can endure impacts of up to 30 Ns from the front and over 40 Ns from the back, although it tends to lean backward when deployed, increasing the risk of falling..

D. Locomotion Policy

We used the same three-phase training procedure to train a locomotion policy that achieves a walking speed of 1.2 m/s. In the final phase, we integrated the World Model Denoiser structure, adding a privileged information decoder to enhance training efficiency.

IV. Basic skills with high dexterity

Robots may experience control failures or unintended falls in complex environments, making autonomous fall recovery appealing. We demonstrate the robot's getting-up motion, which consists of multiple joint-level trajectories. The flexibility of the hip joints is utilized to raise the torso and position the center of mass above the support polygon of the feet.

The robot also can flip over when lying on its back, benefiting from the extensive range of motion in its waist joint. With these two basic motions, we can easily design a controller that enables the robot to stand up from any fallen position.

With the ability to get up, intentional falling behaviors become feasible. We showcase the robot's getting-down motion, essentially the reverse of the getting-up motion. Additionally, we have planned a push-up motion, where the robot first squats down to assume a push-up position and then performs repeated push-ups. Through the combination of flipping over and getting up, the robot can independently execute these motions.

V. Future research on the humanoid soccer game

Booster T1 enables greater computing power in robots, paving the way for more learning-based techniques in RoboCup games, resulting in more human-like and entertaining matches. Our platform, equipped with robust computing capabilities and developer-friendly training frameworks, will soon facilitate:

- Learning-based walking: Implementing stable, fast, and resilient gaits in-game strategies.
- Vision-integrated techniques: These techniques simplify ball handling for RoboCup teams through commands for dribbling, shooting, and saving, all trained in a simulated environment.
- End-to-end game training: Allowing two teams to train directly against each other in simulation or real

environments, with learning managing localization, navigation, and strategy..

VI. Conclusion

This paper introduces Booster T1, a Teen-Sized humanoid robot with hardware details and an RL-based locomotion framework. This standard platform enhances the mobility and intelligence of soccer robots for RoboCup participants while simplifying application development with RL and LLMs. The advantages of the standard platform make team cooperation easier, and a joint-team match will be exciting. An 11Vs11 match of humanoid robots will be organized nearly.

A demonstration of the locomotion and push recovery policies on Booster T1 is shown in the following link: https://youtu.be/b-wKLEJXnIk.

References

[1] H. Kitano and M. Asada, “The RoboCup humanoid challenge as the millennium challenge for advanced robotics,” Adv. Robot., vol. 13, no. 8, pp. 723-737, 2000.

[2] Soccer R C. Humanoid League DRAFT Roadmap from 2022 to 2050[J]. URL: http://humanoid. robocup. org/wp-content/uploads/roadmap_2022. pdf (visited on 2023/04/03), 2023.