


Tech United Eindhoven Team Description 2026

R.M. Beumer, D.M.J. Hameeteman, A.S. Deogan, J.J. Olthuis, A. Anil, W. Beks, W. Bocian, R. van den Bogaert, P.E.J. van Brakel, M. Briegel, D.J.H. Bruijnen, A. Chiscuzzu, S. Doodeman, A. Guillot, D. Joshi, S.T. Kempers, K.A. Knapik, I. Kolodko, D. Kriellaars, A. Krishnaswamy, G.M. Langhout, H.C.T. van de Loo, F.B.F. Schoenmakers, P. Teurlings, and M.J.G. van de Molengraft 

Tech United Eindhoven
De Rondom 70, P.O. Box 513 5600 MB Eindhoven, The Netherlands
techunited@tue.nl
<https://www.techunited.nl/>

Abstract. The Tech United Eindhoven Middle-Size League (MSL) has become world champion eight times, including at RoboCup 2024. Moreover, the team achieved first place in the Technical Challenge and second place in the Scientific Challenge at RoboCup 2024. Since then, the team has made considerable developments in various areas. In this work, we highlight the automatic dataset generation for training vision networks, incorporating algorithms that calibrate the robots automatically, and the design and development of a goalkeeper arm mechanism. This paper describes how these developments will contribute to Tech United’s goal of again becoming world champion at RoboCup 2026.

Keywords: RoboCup Soccer · Middle Size League · AI Vision · Automatic Calibration · Goalkeeper Arm

1 Introduction

Tech United Eindhoven represents Eindhoven University of Technology (TU/e) in the RoboCup competition. The team joined the Middle Size League (MSL) in 2006 and played in 14 finals of the world championship, winning it eight times. The MSL team consists of 5 PhD, 6 MSc and 2 BSc students. Furthermore, we have 4 TU/e staff members and 8 former TU/e students and employees. This paper describes the major scientific improvements of the Tech United soccer robots over the past year and elaborates on some of the main developments for RoboCup 2026.

This paper starts with a description of the fifth generation soccer robots to be used during the RoboCup 2026 competition in Section 2. Then, we present two developments to increase the level of autonomy of our robots. Section 3 explores the methodology for generating synthetic datasets to train convolutional neural network-based object detection algorithms, removing the human labelling steps. Section 4 discusses strides made towards enabling robots to perform self-calibration, reducing the need for human involvement in match preparation. Advancements made towards a goalkeeper arm mechanism are highlighted in Sect. 5. Finally, Sect. 6 presents the conclusion of this paper.

2 Robot Platform

The Tech United soccer robots are called TURTLES, which is an acronym for Tech United Robocup Team: Limited Edition. Their development started in 2005, and through years of experience and numerous improvements they have evolved into the fifth generation TURTLE, shown in Fig. 1.




Fig. 1: Fifth generation TURTLE robots, with the goalkeeper on the left-hand side. (Photo by Bart van Overbeeke)

The software controlling the robots consists of four modules: Vision, World-model, Strategy, and Motion. These modules communicate with each other through a real-time database (RtDB) designed by the CAMBADA team [2]. Further details on the software architecture can be found in [7].

3 Synthetic Data Generation

Due to the increasing progression of AI vision capabilities, more MSL teams are adopting this method for object detection. Convolutional Neural-Networks (CNNs) commonly outperform traditional methods whilst providing robust detection for a wide range of objects [6]. These networks are tuned by feeding them with many annotated images (on the order of thousands for small datasets). For MSL, the model uses annotations of the robots and balls to learn and recognize patterns and features in the labelled images. Currently, teams manually capture and annotate images, which is a long and gruelling process prone to error.

Existing labelled datasets can be used for training, drastically speeding up the annotation process. Methods that employ existing CNN models to label new datasets cannot be used in MSL due to the lack of publically available, robust models capable of detecting MSL robots. Therefore, we investigate synthetic dataset generation to rapidly and automatically generate vast annotated datasets.

3.1 Simulation Environment

Figure 2 shows the environment for the generation of synthetic data. Unreal Engine is used as the simulation environment, LUMA AI is used for the generation of photorealistic 3D Gaussian splat (3DGS) models [5]. 3DGS's provide photorealistic models capable of real-time rendering, enabling rapid generation of large-scale, diverse synthetic datasets.



Fig. 2: Example of the virtual environment, including the field and goals, the lights, the background augmentation, and 3D Gaussian splats of balls and robots.

For the generation of representative synthetic datasets, it is critical to identify the factors that need to closely mimic the real world. To ensure that the network trains on the actual features of the objects, the most important factor is photorealism of the 3D models of the objects. These photorealistic models are created by capturing images of the object, where only a normal smartphone camera is required, and feeding them into LUMA AI, which generates the 3D models within an hour [5].

The main advantage of this method is the ease of introducing domain randomization in the dataset. Changing the lighting conditions, object orientations, lens parameters, and camera orientations are standard features in many simulators. Applying a uniform distribution, with predefined bounds, to these variations provides an automated method of gathering synthetic images. Furthermore, the calculations of the 2D bounding boxes are automated by projecting the 3D coordinates of the objects onto the 2D image plane and capturing the extents in the screen space [1].

Generating one labelled synthetic image using this method takes approximately 1 second, drastically speeding up the process of gathering and annotating diverse datasets for training a robust object detection algorithm.

3.2 Results

Trainings were conducted based on YOLOv8 with 150 epochs, a batch size of 16, and approximately 2500 images with Automatic Mixed Precision (AMP) training. Additional parameters used are the default training settings provided by Ultralytics [4]: an initial and final learning rate of 1×10^{-2} , a momentum of 0.937, and a weight decay of 5×10^{-4} .

Table 1 quantifies the method in match gameplay using conventional metrics, namely, precision, recall, F1-score, and mean average precision (mAP). mAP50 provides a single number metric to summarise the detection performance of the model in multiple classes of objects. Similarly, mAP50-95 is determined across ten intersection-over-union thresholds ranging from 0.5 to 0.95, with increments of 0.05. In robot soccer, precision is more important than recall, as false positives lead to incorrect decisions, such as detecting background objects as the ball.

The mAP50 being around 87% shows that our method provides a balanced and solid approach, providing a good starting point for generating synthetic data for the MSL. The lower recall is due to the limited size of the dataset used, however, in the future, the number of images will be increased significantly which will increase the recall accordingly. Furthermore, using the detection algorithm alongside an estimation and tracking algorithm reduces the system’s dependency on a high recall. Especially since this high precision low recall ratio means that we limit the number of undesirable false positives (ghost obstacles), only strengthening the estimation and tracking algorithms.

Figure 3 shows an example of the synthetic data trained YOLO performance on the validation dataset. This checks how well it performs in a real-world match scenario with motion blur. Since the datasets do not inherently model motion blur, it is possible that this can affect the performance of the network. However, tools such as Roboflow could be used, which use AI techniques to introduce motion blur in static images.

A more detailed explanation of this work is given in [3]. Given these results, our aim is to use this method to detect robots and balls at RoboCup 2026.

4 Automatic Calibration

In RoboCup, the goal is to develop fully autonomous robots capable of playing soccer without human intervention. While robots are able to operate autonomously during matches, a lot of human involvement is required for match preparation. Various aspects of our robots, including ball detection, localization on the field, ball handling, and shooting mechanisms, require periodic calibration. One key issue is the need for constant recalibration, particularly for ball

Table 1: Validating our object detection method using generated datasets on real life match gameplay.

mAP50↑	mAP50-95↑	Precision↑	Recall↑	F1-score↑
0.868	0.7037	0.953	0.759	0.845




Fig. 3: Example of the validation dataset from real life match gameplay, where ‘p1’ represent the robot and ‘b’ represents the ball.

detection, which currently depends on lighting conditions. By implementing AI vision, as discussed in Sect. 3, we can create a more robust ball detection system that adapts to changing lighting without needing recalibration. The other aspects of the robots that require frequent recalibration are more mechanical in nature. For example, to shoot the ball, the robots need to know the physical alignment of the shooting mechanism with respect to the vision system. Collisions with obstacles and other robots, which still frequently occur in a match, lead to slight changes in this alignment. It would therefore be very beneficial if the robots could correct for these changes in the calibration values in real-time during a match. As a first step, we want the robots to perform this real-time calibration outside the match, so that we still have control over the calibration. Then, when this proves to be robust and reliable, we can incorporate the real-time calibration into our match play.

We propose the robotic equivalent of a human warming-up session. During this warm-up, the robots will perform real-time calibration and test the overall system functionality to ensure everything is working properly. In a simple setup, the robots move in a predefined pattern while passing the ball to each other, allowing them to calibrate both their shooting and receiving mechanisms. More advanced warming-up sessions can include shooting at the goal to fine-tune the height and speed of the robots’ shots.

Currently, the real-time calibration of the angle between the vision system and shooting mechanism is implemented. Up until now, calibration has been done by manually performing a shot, taking an image of the ball leaving the shooting mechanism, and adjusting the value of the shooting angle in the software. This procedure is now replaced by an automatic calibration procedure. Whenever the robot passes, it will track the direction in which the ball is moving by analysing the camera data. This is then compared to the location of the intended target. The angle between the two is then used to adjust the calibrated shooting angle. Figure 4 shows a situation in which the intended ball trajectory (red) does not align with the actual shooting angle (yellow).

In a parallel project, we use real-time data analysis to determine whether a recalibration of the robot is necessary. When the robot determines one of its calibration values is no longer correct, for example, after a collision, it will request a substitution to allow it to be fine-tuned in the team technical area. In the future, the combination of real-time calibration possibilities and this validation of the calibration values during play allow for the robot to fine-tune its values when it deems necessary to do so.



Fig. 4: An example of an image taken by the robot after kicking the ball. The red arrow shows the expected direction of the shot, while the yellow arrow indicates the actual direction of the shot.

5 Goalkeeper Arm

The primary objective of our new goalkeeper arm is to extend the robot's effective reach to cover a lateral region 60 cm from each side and 20 cm above the robot while remaining compliant with the RoboCup MSL convex-hull and timing rules. The mechanism must withstand ball impacts up to 22 m/s, be capable of extending to any required point in the designated region within 0.5 s and be retracted within 1 s, and achieve a positioning accuracy of approximately 5 cm.

The mechanical concept centers on a three-stage telescopic extension built from carbon-fiber reinforced polymer tubes to maximize stiffness-to-mass ratio. Translation is actuated using a belt/cable system that enables bi-directional control with minimal inertia on the moving stages. Between the rotational and translational domains, an integrated energy-dissipation element (spring plus damper) was introduced to absorb impact energy and limit peak loads transmitted to the structure. Guidance and bearing choices include roll-wrapped carbon tubes running in polymer plain bearings, and a dual-tube arrangement at the end effector to prevent undesired rotation and increase bending stiffness.




Fig. 5: Overview of the keeper arm, including damping components (red), extension components (blue) and actuation components (green).

Experimental validation showed that the rotational axis met the specified speed and steady-state accuracy targets, and that the translational axis achieved the required steady-state positioning accuracy. Translational deployment speed, however, approached but did not fully meet the 0.5s requirement, and the prototype mass still exceeded the original 5 kg target. Under combined motion the translational axis exhibited degraded tracking and reduced speed attributable to guidance friction and transmission losses, identifying clear mechanical and actuation areas for improvement.

All in all, the rotation-translation architecture with telescopic tubes and a dedicated dissipation axis provides a viable mechanical baseline for high-speed interception within MSL constraints. Coupled ball-arm modeling was essential to size structural elements and specify damping, and the prototype validated the core concept by meeting accuracy and rotational speed targets while highlighting translational speed and mass as the primary development priorities.

6 Conclusion

In summary, this paper outlines the major scientific and technological advancements of the Tech United soccer robots since their previous participation in RoboCup.

The introduction of synthetic datasets for training convolutional neural network-based object detection algorithms has been highlighted, demonstrating their effectiveness in improving performance for robot soccer while significantly accelerating the annotation process. This innovation simplifies the inclusion of additional robots and objects into the system.

Furthermore, progress towards more autonomous calibration was presented, including the implementation of automatically calibrating the robot's heading direction for kicking. This approach reduces the reliance on human intervention and takes important steps toward achieving the vision for 2050, where robots will independently warm up by calibrating their sensors and actuators.

Lastly, the integration of a goalkeeper arm, as well as the development of our new sub-team to participate in the humanoid league (not described in this team description paper, but done in parallel), aligns closely with the long-term objective for 2050: beating the human world champions with a team of humanoid robots.

References

1. 2D Bounding Box for Deep Learning Image Detection (2024), <https://www.fab.com/listings/46a82ff2-8e9b-4a04-9434-7c7c8bebac63>, accessed: 2025-01-22
2. Almeida, L., Santos, F., Facchinetti, T., Pedreiras, P., Silva, V., Lopes, L.S.: Coordinating distributed autonomous agents with a real-time database: The CAMBADA project. In: Computer and Information Sciences - ISCIS 2004. pp. 876–886. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
3. Deogan, A., Beks, W., Teurlings, P., De Vos, K., Van Den Brand, M., Van De Molengraft, R.: Synthetic dataset generation for autonomous mobile robots using 3d gaussian splatting for vision training. In: 2025 IEEE 21st International Conference on Automation Science and Engineering (CASE). pp. 2802–2807 (2025). <https://doi.org/10.1109/CASE58245.2025.11163770>
4. Jocher, G., Qiu, J., Chaurasia, A.: Ultralytics YOLO (Jan 2023), <https://github.com/ultralytics/ultralytics>
5. LumaAI: Gaussian Splats for 3D Visualization (2025), <https://lumalabs.ai/>, accessed: 2025-01-22
6. O'Mahony, N., Campbell, S., Carvalho, A., Harapanahalli, S., Hernandez, G.V., Krpalkova, L., Riordan, D., Walsh, J.: Deep learning vs. traditional computer vision. In: Advances in computer vision: proceedings of the 2019 computer vision conference (CVC), volume 1 1. pp. 128–144. Springer (2020)
7. R. M. Beumer et al.: Tech United Eindhoven Middle Size League Winner 2023. In: RoboCup 2023: Robot World Cup XXVI. pp. 428–439. Springer Nature Switzerland (2024)