ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

Scientific and Technological Challenges in RoboCup

Minoru Asada[1] and Oskar von Stryk[2]

[1]Symbiotic Intelligent Systems Research Center, Institute for Open and Transdisciplinary Research Initiatives, Osaka University, Suita, Osaka 565-0871, Japan; email: asada@otri.osaka-u.ac.jp

[2]Simulation, Systems Optimization, and Robotics Group, Department of Computer Science, Technische Universität Darmstadt, D-64289 Darmstadt, Germany

Annu. Rev. Control Robot. Auton. Syst. 2020. 3:441–71

First published as a Review in Advance on January 13, 2020

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-100719-064806

Copyright © 2020 by Annual Reviews. All rights reserved

ANNUAL REVIEWS CONNECT

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

RoboCup, challenges, soccer, rescue, home, education, industry

Abstract

Since its inception in 1997, RoboCup has developed into a truly unique and long-standing research community advancing robotics and artificial intelligence through various challenges, benchmarks, and test fields. The main purposes of this article are to evaluate the research and development achievements so far and to identify new challenges and related new research issues. Unlike other robot competitions and research conferences, RoboCup eliminates the boundaries between pure research activities and the development of full system designs with hardware and software implementations at a site open to the public. It also creates specific scientific and technological research and development challenges to be addressed. In this article, we provide an overview of RoboCup, including its league structure and related research issues. We also review recent studies across several research categories to show how participants (called RoboCuppers) address the research and development challenges before, during, and after the annual competitions. Among the diversity of research issues, we highlight two unique aspects of the challenges: the platform design of the robots and the game evaluations. Both of these aspects contribute to solving the research and development challenges of RoboCup and verifying the results from a common perspective (i.e., a more objective view). Finally, we provide concluding remarks and discuss future research directions.

1. INTRODUCTION

Since its inception in 1997, by offering a publicly appealing yet formidable challenge (1–3), RoboCup has been a vehicle for promoting highly ambitious research in robotics and AI. One effective way to promote science and engineering research is to set a visionary and challenging long-term goal, and RoboCup was founded with such a long-term goal: by 2050, to have a team of soccer-playing robots defeat the most recent World Cup champion team. This goal was set at a time when humanoid robots were still confined largely to science fiction, as the Honda P2 humanoid robot was unveiled only in December 1996. Besides huge challenges for technology, RoboCup's vision also raises philosophical and societal questions (4).

This challenge has been successively expanded to address societal challenges by including major leagues for rescue robots, robots that perform services for humans at home, and robots that perform manufacturing tasks. The junior leagues, which target children in primary and secondary school as well as undergraduates under 19 years old, comprise robotic soccer, rescue robots, and creative on-stage performances by robots and humans. Generally, building teams of robots that perform services and operate in environments with a large amount of uncertainty (such as soccer games and rescue operations) can have significant social and economic impact, and reaching the specific 2050 goal would certainly be a major achievement in the science and engineering fields of robotics and AI.

Figure 1 illustrates how the number of RoboCup leagues has expanded since 1997. The first RoboCup had three soccer leagues [the Simulation League, Small Size League (SSL), and Middle Size League (MSL)] and has since expanded to five domains, each comprising several leagues: RoboCupSoccer, RoboCupRescue, RoboCupJunior, RoboCup@Home, and RoboCupIndustrial. In the following sections, we briefly explain the main research issues and their variations in each league. The number of participating teams increased rapidly in the first 10 years, but owing to

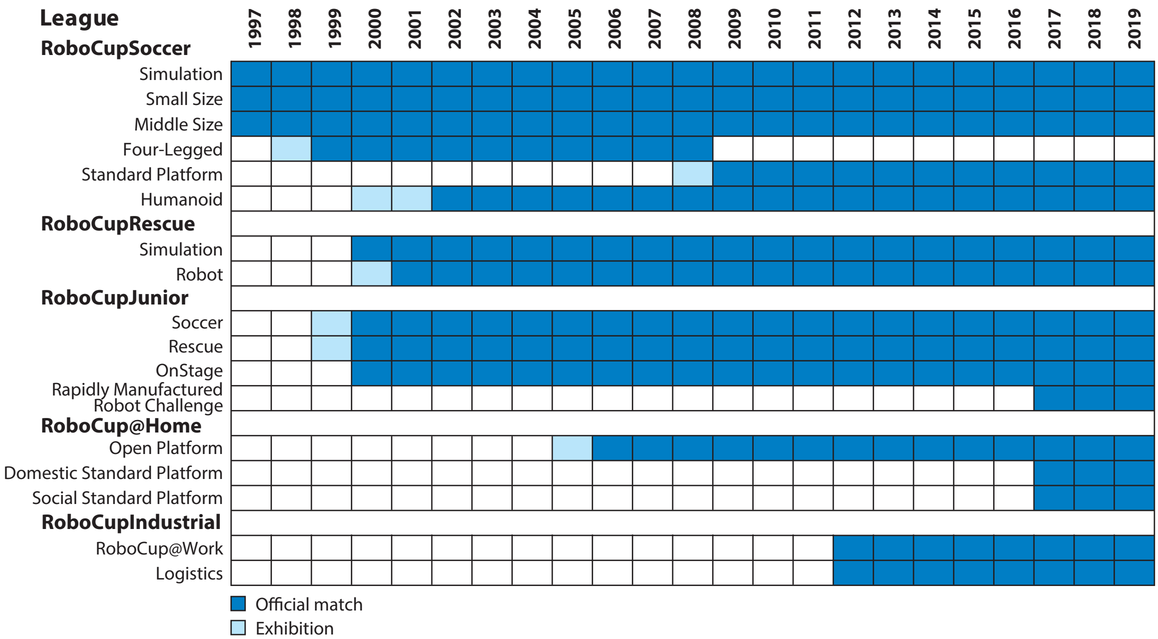

Figure 1

Expansion of RoboCup leagues since 1997. All leagues are currently active except for the RoboCupSoccer Four-Legged League, which was replaced in 2009 by the Standard Platform League.

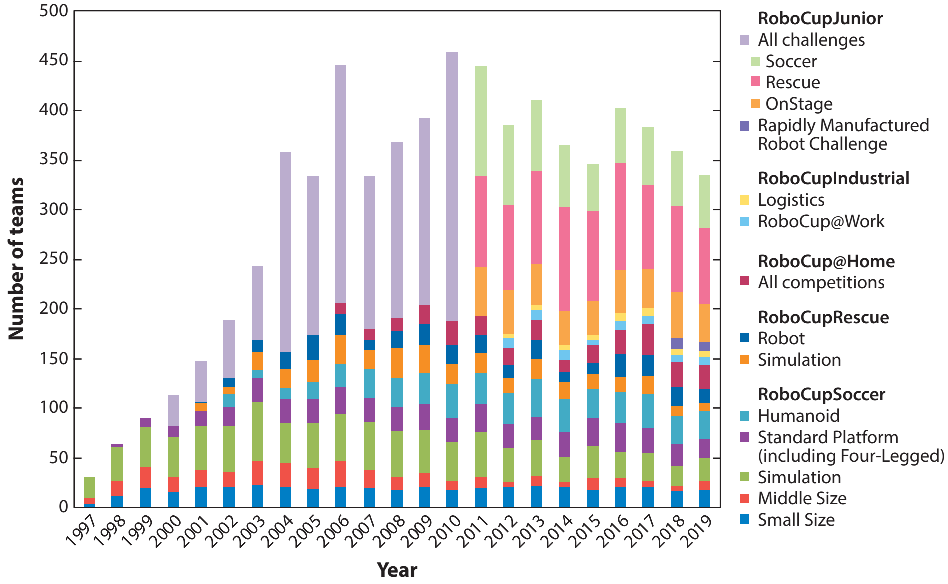

Figure 2

Changes in the total number of teams across RoboCup leagues since 1997. Because league-specific breakdowns for RoboCupJunior are not available for 2000–2010, those years show total numbers of RoboCupJunior teams for all challenges.

the limited space and time, the number of teams in the international global event is now limited to approximately 350–450 in the major and junior leagues combined (see **Figure 2**). There are approximately 1,000 major league teams and 10,000 junior league teams worldwide. There are also many regional (e.g., the Japan, German, Iran, and Portuguese Opens) and supraregional (e.g., the Asia-Pacific Open) RoboCup events that offer further opportunities for active participation. The current structure of the RoboCup leagues, their current rules and committees, and information on how to participate in RoboCup are available at the RoboCup website (**http://www.robocup.org**).

The remainder of the article is structured as follows. In Section 2, we explain why RoboCup is unique as a research community and how the league structure has expanded. Next, we survey the research issues based on the broad research categories and topics. We then review a number of current studies to show how the participants (called RoboCuppers) identify and address related challenges. Finally, we provide concluding remarks and map future research directions.

2. WHY IS ROBOCUP UNIQUE?

To reflect the basic policy mentioned in Section 1, RoboCup has removed several types of boundaries:

1. The boundary between competition organizers and participants: The technical committee members of each league include members elected from among team representatives, and the executive committee members are elected from among the former technical committee members. Thus, the boundary between the participants and organizers is naturally blurred, and they participate equally in discussing and designing new leagues and challenges.
2. The boundary between short- and long-term achievements: Unlike many robotics competitions, which focus exclusively on significant short-term development achievements, RoboCup focuses on sustainable, long-term progress toward its ultimate goal.

3. The boundary between academia and industry: Intelligent robotics is intrinsically interdisciplinary, encompassing many fields in science and technology, and many RoboCup studies similarly cover a wide range of disciplines rather than very specific issues in narrower areas. Furthermore, RoboCup has an active relationship with the robotics industry in both directions: New companies are born from the RoboCup community, and companies hold their own challenges or support leagues and/or teams in RoboCup. For example, the leader of the 2002 SSL champion team, Professor Raffaello D'Andrea, and his colleagues started a company for mobile robotic fulfillment systems called Kiva Systems in 2002. Kiva Systems produced mobile robot systems for carrying shelves in warehouses based on technology fostered by RoboCup. Amazon acquired the company in 2012 and renamed it Amazon Robotics (**https://www.amazonrobotics.com**) in 2015, and in 2016 and 2017, Amazon Robotics held its own RoboCup competitions, the Amazon Picking Challenge (2016) and Amazon Robotics Challenge (2017). In the other direction, Sony's AIBO robot was used as the standard platform in the Four-Legged League until 2008, when RoboCup started the Standard Platform League (SPL) and selected the Aldebaran NAO humanoid robot as the new standard platform following an open call for tenders; this was the first major application of the NAO robot. SoftBank acquired Aldebaran Robotics in 2015 and renamed it SoftBank Robotics Europe (SoftBank Robotics having been established in 2014). In 2014, a new service robot, Pepper, was introduced, and in 2017 it was selected as the social standard platform for the RoboCup@Home domain following another open call for tenders. The Toyota Human Support Robot (HSR) was also selected as a domestic standard platform in the RoboCup@Home domain following the same call. The Quince robot for disaster response developed by Tohoku University has been evaluated in the RoboCupRescue Robot League in various iterations, and it was finally deployed at the Fukushima Daiichi nuclear power plant in the aftermath of the massive earthquake and tsunami that hit eastern Japan in 2011, where it was used for inspection missions in highly contaminated areas (5).
4. The boundaries between nations: Unlike other competitions, such as the Olympics and the FIFA World Cup championship, RoboCup is not focused on competition between nations. Rather, it encourages international collaborations of joint teams and many research studies by authors of different nationalities. These collaborations have been supported by strong human networks fostered in the RoboCup community, and RoboCupJunior enhances these networks even more. Every year, participants from approximately 40–45 different nations and regions participate in the annual RoboCup event with a friendly, cooperative, and enthusiastic spirit.
5. The boundary between university-level research and project-oriented STEM education for primary and secondary school children: RoboCupJunior and its links to the goals of the major leagues have effectively eliminated this boundary.

As demonstrated by the above activities and achievements, RoboCup is advancing autonomous and intelligent robotics. It was the first organization to introduce visionary competitions for intelligent robots, including multiagent teams of autonomous robots, and is the oldest active robotics competition; for these reasons, it is referred to as "the mother of all competitions." RoboCup pioneered the idea of benchmarking robotic systems through competitions, including on the functional level, addressed by technical challenges, and on the system and mission levels, addressed by games. Moreover, RoboCup was the first organization to introduce competitions for intelligent rescue, home, and industrial robots, otherwise known as the forerunners of "Industry 4.0" and "Society 5.0," in efforts predating these popular terms.

3. LEAGUE STRUCTURE

Through its board of trustees and executive committee, RoboCup has expanded the number of leagues since 1997, as shown in **Figure 1**. In the RoboCupSoccer domain (6), vision systems are the main external sensors used for object detection (the ball, goals, teammates, opponents, field lines, etc.) and localization. The SSL uses a global open-source vision system called SSL-Vision (7) that utilizes one or more standard cameras on the ceiling (see **Supplemental Figure 1**). Using this system relaxes the requirements of the robots' usually rather limited onboard perception and capabilities, enabling a highly dynamic game with six robots per team. The MSL uses individual onboard camera systems. The SSL-Vision system enables hybrid centralized/distributed control of many robots by utilizing absolute position information for each robot and the ball, whereas the MSL onboard vision system makes it necessary to control each robot individually using distributed control with synchronizing mechanisms. The MSL initially used a standard camera for onboard vision, but to increase the visual field and reduce uncertainty, an omnidirectional vision system is popularly used as the de facto standard.

Supplemental Material >

In the SSL and MSL, movement around the field is enabled by wheel-based locomotion. A differential drive was originally used in which each robot was propelled by two wheels and changed direction by varying the rotation speed of the two wheels (a nonholonomic-type system). However, the current practice is to use an omnidirectional vehicle as the de facto standard for quick movement in any direction. The Four-Legged League started in 1998, and the Sony AIBO was the first standard platform used in RoboCup. Bipedal locomotion using a human-like body plan is the focus of the Humanoid League (8), which has three classes differentiated by the size and number of players: KidSize, TeenSize, and AdultSize. The Humanoid League originally used multi- and omnidirectional vision systems but allowed only external and internal sensors that have a rough equivalent in human senses, including a human-like field of view. Every few years, the real robot soccer leagues have increased the field size and/or number of players and made the setup of the game environment more realistic, making it more difficult to effectively tackle the numerous challenges. **Figure 3** shows the playing fields for the SSL, MSL, and Simulation League at the first RoboCup in 1997 and the latest RoboCup in 2019. The SSL field was initially a ping-pong table (152.5 cm × 274.0 cm) and in 2019 was 900 cm × 1,200 cm (approximately 26 times larger). Similarly, the MSL field size has increased from a 3 × 3 set of ping-pong tables (totaling 457.5 cm × 822.0 cm) in 1997 to 14 m × 22 m in 2019 (more than 8 times larger).

The 2-D Simulation League uses a local vision system with a limited visual angle and an omnidirectional locomotion system. It also uses a public auditory system, meaning that players can communicate with their teammates and that their opponents can hear any conversations. After the Sony AIBO was retired as the standard platform in the Four-Legged League, the NAO humanoid robot from Aldebaran Robotics was selected as the new standard platform for the SPL, and the 3-D Simulation League uses a digital NAO model.

The RoboCupRescue domain allows teleoperation by humans using the robots' onboard sensory information; although full autonomy is ideal in robotics, it is particularly challenging in the highly unstructured environment of a disaster site, and the environments vary across missions. The task of rescue robots is rather different from those of intelligent robots in structured and undamaged domains, such as industries or private homes. Therefore, fully autonomous rescue robots are especially problematic and are still far from practical application in critical missions such as real rescue operations (e.g., saving a disaster victim within 72 hours). Therefore, autonomous capabilities are expected to be introduced as assistance functions supporting human operators in their tasks. The maneuverability of these mobile robots in highly unstructured (severely damaged) environments is a key issue, and the teams have used and developed various tracked-wheel-type mobile robots.

Figure 3

The fields for the Small Size League (*left*), Middle Size League (*center*), and Simulation League (*right*) at the first RoboCup in 1997 (*top*) and at the latest RoboCup in 2019 (*bottom*). Left and center panels in the top row adapted from Reference 9.

The RoboCup@Home domain aims to develop intelligent service and assistive robot technologies with high relevance for future personal domestic applications. Human–robot interaction and cooperation are among the core issues in the development of these robots. It is also necessary to hone their object recognition skills, as they are required to recognize many types of objects under natural light conditions in daily life. Skills such as grasping and manipulating objects and mapping and navigating a changing environment inside a building, such as a living room, office, or supermarket, are particularly vital. Communication with humans is necessary to respond to human instructions. As such, speech recognition, understanding, and responses to humans are required. Unlike the RoboCupSoccer domain, RoboCup@Home is not based on an adversarial environment; instead, it thrives on collaboration between robots and humans.

The RoboCupIndustrial domain comprises two leagues: RoboCup@Work and the RoboCup Logistics League (RCLL). Both are oriented toward industry. RoboCup@Work focuses on mobile robots with manipulators that cooperate with and assist humans in a futuristic industrial production environment, while the RCLL consists of multiple mobile robots cooperatively planning, executing, and optimizing the material flow and product delivery according to dynamic orders in a smart factory environment. The standard platform of the RCLL is Robotino, a mobile robot from Festo. Although both leagues have many research issues in common with other domains and leagues, they require industrial qualification for practical applications.

The competitions in the RoboCupJunior domain focus on project-oriented education and edutainment and mainly involve robots participating in two activities inspired by the major leagues, soccer and rescue. The competition unique to RoboCupJunior is OnStage, wherein teams develop a creative stage performance and compete against one other using autonomous robots that they have designed, built, and programmed. The three competitions are designed to enable children and teens to simultaneously learn several subjects (such as physics, mechanics, electrotechnics, and electronics) through computer programming without any formal separation of the subjects (as would be the case in, e.g., a typical university curriculum). Furthermore, as they execute their project, the participants naturally learn social teamwork.

Table 1 summarizes the differences among the leagues. In subsequent sections, we discuss research achievements with reference to this table.

4. A TREASURE TROVE FOR A RICH DIVERSITY OF RESEARCH ISSUES

Many research issues in RoboCup are structured in terms of categories and specific topics (see **Table 2**), with an eye toward the goals of the four major domains. Although this classification is not strict, it provides a guide, as many research endeavors cover different topics and categories owing to their interdisciplinary nature. For example, although 3-D perception and robot kinematics and dynamics are separate challenges, they are combined and studied under the category of sensory-motor control. Furthermore, these combined challenges are supported by infrastructure (robot hardware and software) challenges for the teams participating in the real robot leagues, especially the ones that do not use a standard platform. It should be noted that the complexity of the robot hardware and corresponding software increases from the SSL through the MSL and on to the RoboCupRescue domain and Humanoid League. Progress in overcoming these challenges is verified through competitions that are open to the public (see **Supplemental Figure 2**).

Supplemental Material >

Among the seven categories in **Table 2**, two categories—robot hardware and software, and applications and benchmarking—are special, as many RoboCuppers can be inhibited by hardware-related issues that increase the time and effort required to design, assemble, and maintain the robots. This is the main reason why RoboCup pushed standard platforms in several domains: Shared platforms make evaluations regarding the capabilities of software modules and functionalities more objective, which is important in both research and applications.

The achievements at RoboCup can have a broad impact; in addition to providing great education and training for participants in RoboCup, the work significantly impacts technology transfer and development. Several RoboCup achievements have been presented at the annual RoboCup symposiums held directly after the competitions and published in the corresponding proceedings (10–32), as well as presented at regular international robotics and AI conferences. In the following sections, we mainly review recent activities described in the symposium papers of the last five years (28–32), along with a small number of earlier studies.

5. ROBOT HARDWARE AND SOFTWARE

Robot hardware and software are central in any study of intelligent and autonomous robotics. RoboCup in particular necessitates various types of robot platforms, depending on the specifications of the different leagues. Below, we review first the hardware platforms and then the software platforms.

5.1. Hardware Platforms

Because no hardware platforms were available when RoboCup began, all the robots were originally either built from scratch or based on modifications to commercially available parts. The MSL of the first RoboCup had five teams, and the two cochampions—Trackies, from Osaka University (**Figure 4*a***), and Dreamteam, from the University of Southern California Information Sciences Institute (**Figure 4*b***)—coincidentally modified the same Japanese radio-controlled toy cars, which implies that few suitable platforms were available. Two other teams used omnidirectional

Table 1 Key features of different RoboCup leagues

Domain or league	Number of robots	External sensors	Objects to be recognized	Research topics	Motion mechanism
RoboCupSoccer SSL	5–8	Global and fixed camera(s) shared by both teams	Ball, goals, field lines, teammates, and opponents	Real-time detection and tracking of objects, motion control, and teamwork	Wheel-based, omnidirectional locomotion with kicking device
RoboCupSoccer MSL	5	Onboard individual cameras (mainly horizontal motions) and omnidirectional vision systems	Ball, goals, field lines, teammates, and opponents	Real-time detection and tracking of objects, integration of individual robots' visual information into a global map, motion control, and teamwork	Wheel-based, omnidirectional locomotion with kicking device
RoboCupSoccer SPL and Humanoid League	5 (SPL, Humanoid League KidSize), 2 (Humanoid League TeenSize and AdultSize)	Onboard individual cameras (active vision through 3-D motion) with field of view limited to 180°	Ball, goals, field lines, teammates, and opponents	Real-time detection and tracking of objects, integration of individual robots' visual information into a global map, localization with a limited field of view under significant disturbances, motion control, complex robot behavior control, and team coordination and cooperation	Posturally stable human-like bipedal locomotion with various types of whole-body motion
RoboCup@Home	1	Onboard multiple cameras, common 2-D and 3-D sensors (e.g., lidar and RGB-D), and microphones	Humans, including their intentions and status; many daily-life objects; and the environment and its current status	Recognition of many objects and subjects and their localization and manipulation, communication, understanding and fulfilling tasks specified by humans on a human-like level of description	Wheel-based locomotion with arm and hand
RoboCupRescue	1–2	Onboard multiple color, 3-D, and thermal cameras; lidar; and microphones	Status of environment, including victims and damage found autonomously or with support from human operators	Moving from teleoperation to more autonomy in tasks such as mobility, navigation, and manipulation	Mainly tracked wheels for locomotion with arm and hand
RoboCupIndustrial	1 (RoboCup@Work), 3 (RCLL)	Onboard color and 3-D cameras and lidar (RoboCup@Work); camera, lidar, and multiple infrared distance sensors (RCLL)	Industrial parts (RoboCup@Work); status of environment in exploration and production phases (to discover and utilize production machines) and status of other robots in the same environment, including other teams (RCLL)	Recognition and autonomous manipulation of industrial parts, transportation, and navigation (RoboCup@Work); production planning and execution with a mobile multiagent team in a smart-factory environment (RCLL)	Wheel-based locomotion with arm and hand

Abbreviations: MSL, Middle Size League; RCLL, RoboCup Logistics League; RGB-D, red, green, and blue plus depth; SPL, Standard Platform League; SSL, Small Size League.

Table 2 Selected research categories and topics in RoboCup

Category	Topics
Robot hardware and software	Mobile robotics, humanoid robotics, sensors and actuators, embedded and mobile devices, robot construction and new materials, robotic system integration, robot software architectures, robot programming environments and languages, real-time and concurrent programming, robot simulators
Perception and action	3-D perception, distributed sensor integration, sensor noise filtering, real-time image processing and pattern recognition, motion and sensor models, sensory-motor control, robot kinematics and dynamics, high-dimensional motion control
Robot cognition and learning	World modeling and knowledge representation; learning from demonstration and imitation; localization, navigation, and mapping; planning and reasoning; decision-making under uncertainty; neural systems and deep learning; complex motor skill acquisition; reinforcement learning and optimization; motion and sensor model learning
Human–robot interaction	Robot social intelligence; fluency of interaction; speech synthesis and natural language generation; natural language recognition; explainable robot behaviors; emotion recognition and reaction; understanding of human intent and behavior; safety, security, and dependability; enabling humans to predict robot behavior
Multirobot systems	Team coordination methods, communication protocols, learning and adaptive systems, teamwork and heterogeneous agents, dynamic resource allocation, adjustable autonomy
Education and edutainment	Robotics and artificial intelligence education, educational robotics, robot kits and programming tools, robotic entertainment
Applications and benchmarking	Search-and-rescue robots; robot surveillance; service and social robots; robots at home, at work, and in public spaces; robots in the real world; performance metrics; human–robot interaction

locomotion systems with different mechanisms: One used omni-wheels (**Figure 4*c***) and one used rolling spheres (**Figure 4*d***). The fifth team used the Pioneer I robot, a nearly complete platform consisting of commercially available parts (**Figure 4*e***). As mentioned above, omnidirectional movements are suitable for quick motion in any directions. Although omni-wheels are commercially available, many teams constructed their own omni-wheels to increase the robots' speed and stability.

Tech United Eindhoven recently demonstrated that the conventional triangular omni-wheel system could not deliver all of the torque from the motors in the desired movement. Furthermore, a high forward acceleration may cause the front wheels to slip, thus preventing the robot from applying torque from the motors to the field. The team developed an eight-wheeled platform to resolve these drawbacks. A key issue is how to resolve an overactuated system with four or more wheels. The eight-wheeled platform (see **Supplemental Figure 3**) has three degrees of freedom and is five times overactuated. Therefore, to allow the five internal movements, each

Supplemental Material >

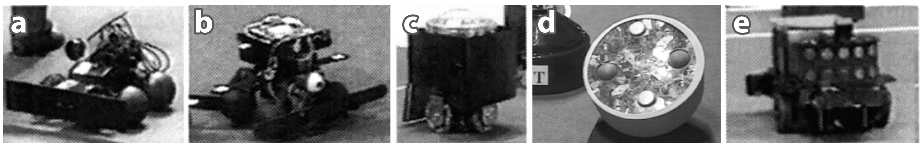

Figure 4

The five platforms used in the Middle Size League at the first RoboCup in 1997. (*a*,*b*) Modified radio-controlled toy cars used by the cochampion teams Trackies (panel *a*) and Dreamteam (panel *b*). (*c*) A robot with omni-wheels for locomotion. (*d*) A robot with rolling spheres for locomotion. (*e*) A Pioneer I robot consisting of commercially available parts. Figure adapted from Reference 9.

of the wheel combinations is suspended, with a rotation point below the ground, and the back wheels are suspended over an axle hinge. Because of this mechanism, the wheels are always in contact with the ground to transfer the torque from the motors to the ground (33). In addition to the locomotion mechanisms, the kicking devices are also custom designed and improved (34–36).

In the SSL, omnidirectional motion control in combination with kicking devices is very popular. One of the SSL teams, OP-AmP, developed a new kicking mechanism with a multiple-angle kicking device using a Geneva drive mechanism to generate straight and diagonal shots in five directions, allowing curved shots by combining the straight kick and backspin behavior (37, 38) (see **Supplemental Figure 4**).

Supplemental Material >

In the SPL, because teams are not permitted to modify the hardware of the standard platform, they are able to focus more on software programming and less on the challenges and damage faced by the robots themselves (which are also not fully avoidable with standard platforms due to the strenuous nature of the competitions). By contrast, in the Humanoid League, many teams have built their own humanoid robots to meet the challenging specifications for a human-like body plan and senses. Compared with the NAO robot in the SPL, robots in the Humanoid League have a much smaller relative foot size (due to the higher ratio of center of mass to allowed foot size) in order to foster research on dynamic humanoid motion and postural stability. Furthermore, the NAO has a non-human-like second camera sensor in its chin that allows it to see a nearby ball on the ground without having to bow its head completely, as in the Humanoid League. Therefore, although the teams in the Humanoid League may prefer to spend more time on behavior generation and other software-based functionalities, they need to spend significant time on designing, building, and maintaining their robots, and a number of humanoid robot platforms (some of which are commercially available) have been developed.

For the KidSize class (height of 40–90 cm), the Dynamic Anthropomorphic Robot with Intelligence–Open Platform (DARwIn-OP) was manufactured by the Korean robot manufacturer Robotis and developed in collaboration with Virginia Tech, Purdue University, and the University of Pennsylvania. DARwIn-OP has 20 degrees of freedom, each controlled by a Dynamixel MX-28T servomotor, and has been useful as a form of standard platform for several years, although recent rule changes that introduced artificial grass on the field have diminished its usefulness. The MX-28T has a stall torque of 24 kgf·cm (at 12 V and 1.5 A) and a 360° range of motion (39), and many teams use Dynamixel servomotors because they are compact, lightweight, easily moduled, and powerful. Fabre et al. (40) proposed an open-source alternative firmware for Dynamixel servomotors; they compared the proposed control strategy with the default strategy and observed significant improvements in terms of accuracy, delay, and repeatability. Bestmann et al. (41) presented a new multibus solution that enables the typical humanoid robots used in RoboCup to have a control-loop frequency of more than 1 kHz, and they also incorporated solutions to integrate sensors into this bus with high update rates. After the success of DARwIn-OP, Schwarz et al. (42) introduced NimbRo-OP as an open platform for the larger classes, such as the TeenSize (80–140 cm height) and AdultSize (130–180 cm) classes in the Humanoid League. NimbRo-OP has a wide-angle camera, ample computing power, and sufficient torque to enable full-body motions, such as dynamic bipedal locomotion, kicking, and getting up after a fall. It is designed to be easily manufactured, assembled, repaired, and modified. Team NimbRo, which developed the robot, has won the championship several times (e.g., 43), including in 2019 (see **Supplemental Figure 5**).

In the RoboCup@Home domain, two standard platforms were introduced in 2017 as a result of an open call for tenders: the Toyota HSR for the new Domestic Standard Platform League, and Pepper from SoftBank Robotics for the new Social Standard Platform League. As in the RoboCupSoccer SPL, the use of these standard platforms avoids the need for teams to spend time building and maintaining their own robots at the expense of other research. The Toyota

HSR was custom developed for RoboCup@Home based on joint research between Toyota and participating teams. It is equipped with many sensors, actuators, and other devices necessary for a wide range of useful tasks (44) (see **Supplemental Figure 6**).

Supplemental Material >

5.2. Software Platforms

Generally, computer simulations are very powerful tools in scientific and engineering studies. In the case of RoboCup, computer simulations are used in several ways:

1. Game environments: 2-D or 3-D simulation leagues are typically conducted using simulated competition environments.
2. Real-time robot simulations: Depending on the research issues, some real robot experiments are difficult or almost impossible to implement or require very large amounts of time and effort. In such cases, various computer simulations of real robots have a long tradition in RoboCup for verifying and/or improving computational methods for perception, action, and planning through suitable digital twins of robots and their interaction with the environment.
3. Proposals for new environments connecting real and virtual worlds.
4. Tools for developing of behavior programming and data analysis.

In relation to the above, several systems and tools have been developed in RoboCup to assist the teams. The Robot Operating System (ROS) from the Open Source Robotics Foundation (**http://osrfoundation.org**) has found widespread use in RoboCup. In addition to other projects, this foundation supports the development, distribution, and adoption of open-source software for use in robotics research, education, and product development.

RoboCuppers have made many open-source contributions to the ROS ecosystem. Moreover, the technical lead for the new, even more capable ROS 2 is Dirk Thomas, a RoboCup graduate who has been active for many years in championship-winning teams in the Four-Legged and Humanoid Leagues. Scheunemann & van Dijk (45) distributed ROS 2 packages to RoboCup teams with benchmarks to show that ROS 2 is a promising candidate for a common framework among leagues. Thielke & Hasselbring (46) proposed a C++ library that compiles neural network models at runtime into machine code that performs inferences. In their experiments on the NAO V6 platform, the library significantly outperformed existing implementations in small networks but was inferior in large networks. Mitrevski & Ploger (47) presented a small Python library for enabling the specification, configuration, and dynamic creation of state machines using a minimal domain-specific language. They demonstrated its validity in scenario definition in contexts such as the RoboCup@Home competition.

Data analysis is an important topic and is expected to be applicable in RoboCup, whose numerous games translate into a massive amount of data. Mellmann et al. (48) presented a system for automatically recording synchronized videos of RoboCup games in the SPL and an application for exploring and annotating large sets of RoboCup-related data. In addition, they provided data sets collected during the 2018 competitions and an algorithm for visually detecting and tracking robots in the RoboCup videos. **Figure 5*a*** provides an overview of the data flow of the implemented data-processing ecosystem, and **Figure 5*b*** shows a sample session for the annotation of kick events regarding their quality in the synchronized video and log data. **Figure 5*c*** shows detected robots in a video recorded using a GoPro camera at the European Open 2016 in Eindhoven, illustrating that object detection and localization are among the most essential tasks in RoboCup. Fiedler et al. (50) developed a tool called ImageTagger that facilitates creating and sharing labeled training data sets for object recognition. ImageTagger is more open and more user-friendly than the existing labeling tools (see **Supplemental Figure 7**).

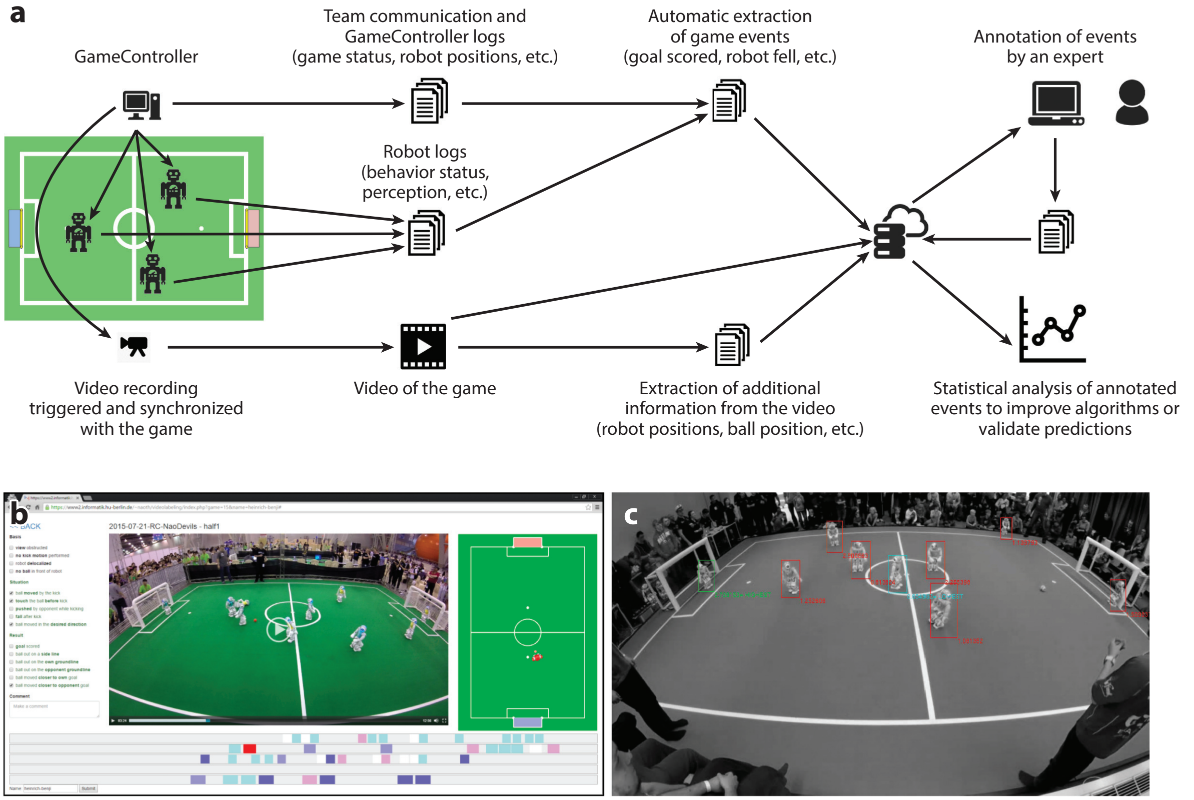

Figure 5

(*a*) Overview of the data flow of the implemented data-processing ecosystem. (*b*) Example of analysis of the quality of kick events. The main components are timelines, with kick events represented by colored buttons (*bottom*), a visualization of the robot's state (position on the field and perceived ball) (*right side*), and an evaluation panel for assigning labels to the events (*left side*). (*c*) An example of detected NAO robots in a video recorded using a GoPro camera at the European Open 2016 in Eindhoven. The colored boxes illustrate the varying confidence levels of the detected robots (49). Figure adapted from Reference 48.

The recent progress in AI technologies, especially deep learning methods, has engendered tasks requiring a large amount of data. Hess et al. (51) proposed a framework for stochastic scene generation and rendering and the automatic creation of semantically annotated ground-truth masks. They evaluated multiple neural network architectures with varying depths and representational capacities and their corresponding runtimes on the current NAO-H25 hardware and provided the required sample training data. Visser et al. (52) proposed a tutorial course that demonstrates how the various tasks can be tackled using the AI and machine learning algorithms available in the MATLAB Statistics and Machine Learning toolbox, and their course works as a toolbox. Van Dijk & Scheunemann (53) proposed a system that processes full VGA images in real time on a low-power mobile processor. Gholami et al. (54) developed another multibody simulation system for a humanoid robot based on MATLAB/Simulink and Simscape software. The system can be used for purposes such as designing control systems and enhancing the stability of robots.

In addition to the useful tools above, new environments have been proposed. Inamura & Mizuchi (55) proposed a novel software platform for statistically evaluating human–robot interaction in competitions. With the help of cloud computing and an immersive virtual-reality system, cognitive and social human–robot interaction can be conducted and measured as objective data in a virtual-reality environment (see **Supplemental Figure 8**). Takami et al. (56) proposed an environment that integrates an agent-development framework and an experiment-management system to support researchers.

Supplemental Material >

Opening the source codes in the RoboCupSoccer Simulation League has played an important role in improving the codes and encouraging new teams to join the league. Code sharing has become much easier and richer in content. Recent examples are the 3-D simulation base code (57), the 2-D simulation base code (58), and an open-source ROS vision pipeline (59). Sharing open-source codes has helped ensure even progress across all leagues [e.g., Hector SLAM (60) and several standard software modules (61) in the RoboCupRescue Robot League].

Introducing referee robots similar to human referees on the field has been one of the significant challenges in RoboCup. Separately from physical referee robots, different types of automated referee software have been developed in the SSL and MSL. The ssl-autonomous-refbox system was an early attempt in the SSL (62); later, Zhu et al. (63) developed AutoRef for monitoring SSL games and detecting rule infringements using data from the global SSL-Vision system. In the MSL, Schoenmakers et al. (64) from Tech United Eindhoven elaborated on the initial steps toward realizing an autonomous MSL referee. As a first step, they implemented an automated referee system that makes decisions based solely on the positioning of the ball and players; the most recent version is available as RefBox (65).

6. PERCEPTION AND ACTION

Perception and action are the most essential research areas in robotics. Some research issues in this category are common across leagues to a certain extent, and some leagues have their own unique issues (see **Table 1**).

6.1. Vision

The global vision system in the SSL has its own vision server (7) that provides low-level visual information, such as the location of the ball and the locations and directions of the robots (teammates and opponents), along with certainty values showing the reliability of the location information. The approach to processing such information depends on the strategy of each team (e.g., 66).

Because teams in the MSL typically use omnidirectional vision systems, a global map is easily obtained by integrating the local maps of teammates and utilizing the information on the locations and directions of the teammates. This can be done without the perspective transformation necessary when using normal perspective cameras (see **Supplemental Figure 9**).

At the inception of RoboCup, the ball was painted red and the goals were painted blue and yellow to simplify image processing and enable quick reactions and movements. One part of reaching the ultimate goal of RoboCup is to make the environment as close as possible to a real human soccer field, and current leagues use a standard black and white soccer ball, which presents a much greater challenge than a colored ball. Modern deep learning techniques can enhance visual recognition tasks even under the constraints of computational resources and quick processing, and many RoboCuppers have applied such techniques by drawing on methods inherent in their task domains.

Speck et al. (67) proposed a neural approach using a convolutional neural network (CNN) to localize the ball in various scenes (see **Supplemental Figure 10**). In the case of a black and white ball, Menashe et al. (68) evaluated and applied a series of heuristic region-of-interest identification techniques and supervised machine learning methods for detecting a ball with high reliability without any prior knowledge of the ball position. Utilizing only black and white images without any color information, Leiva et al. (69) applied a CNN approach to detect a ball and humanoid robots with the following (high) ratios: a robot detection rate of 94.90%, a ball detection rate of 97.10%, a completely perceived orientation rate of 99.88% when the robot under observation was static, and a completely perceived orientation rate of 95.52% when the robot was in motion. However, ball detection is complicated because there are many variations of ball images. To cope with this issue, Teimouri et al. (70) applied CNNs to accurately detect a ball based on an iterative method that employs an efficient integral image-based feature. The features were then fed to a lightweight CNN to finalize the bounding box of the ball, with a resultant detection accuracy of 97.17%. Ball tracking is also a necessary part of the game. Utilizing the temporal information in the CNN, Kukleva et al. (71) presented a system that uses spatiotemporal correlation to efficiently detect and track a soccer ball based on its trajectory. Felbinger et al. (72) designed a CNN for ball detection using a genetic approach that optimized network hyperparameters, providing a cost-effective inference on the NAO with a limited amount of training data.

One constraint on computational resources is the general challenge surrounding the real-time processing of perceptions and actions. In particular, the onboard resources are limited in the SPL, and various teams have therefore devised ways to address this issue. Cruz et al. (73) analyzed the general problem of using CNNs and proposed general design guidelines for their use. Houliston & Chalup (74) proposed an enhancement of CNNs for object detection in resource-constrained robotics through a geometric input transformation called visual mesh. According to the results, the execution time achieved by their method demonstrates outstanding accuracy while being 16 times faster than the fastest competitor tested. Szemenyei & Estivill-Castro (75) proposed an end-to-end neural network solution for scene understanding in robot soccer using two CNNs: one that performs semantic segmentation on an image and another that propagates class labels between consecutive frames. They utilized synthetic data sets and provided RoboDNN, a C++ neural network library. They also extended their study to a system called ROBO, which outperformed Tiny You Only Look Once (Tiny YOLO)—a computer vision system capable of detecting a wide variety of objects in a single image (**https://pjreddie.com/darknet/yolo**)—in terms of both speed and accuracy (76) (see **Supplemental Figure 11**). Poppinga & Laue (77) developed a real-time object detection method for NAO robots called Just Enough Time Net (JET-Net), a model frame for efficiently detecting objects based on CNNs. They reused the learned features to obtain more information from simulation data, a method called simulation transfer learning.

Supplemental Material >

As in the RoboCupSoccer domain, the robots in the RoboCup@Home competitions need to recognize numerous everyday objects. However, in contrast to soccer, transparent objects are common in domestic environments and are particularly difficult to recognize. Hagg et al. (78) proposed a method for recognizing transparent objects using combinations of four modalities: 2-D shapes, 3-D geometry, transparency, and specular reflection. **Figure 6*a*** shows a test set consisting of different types of objects (diffuse, composite, and transparent), and **Figure 6*b*** shows two graphs indicating the performance of the proposed method: one for the entire set of objects, and one for a reduced set with no transparent objects. Using the transparency modality significantly contributes to an increase in the recognition rate even when the set does not include transparent objects.

YOLO and its extensions and modifications are employed for recognizing various types of objects in the RoboCup@Home domain. Reyes et al. (79) proposed a method based on a deep neural network running on a backpack containing a Jetson TK1 card and a battery for near-real-time object recognition for the Pepper robot. Consequently, Pepper could robustly detect and recognize objects in 320 × 320–pixel images at approximately 5 frames per second. Loncomilla & Ruiz-del-Solar (80) proposed YoloSPoC, a method for recognizing particular object instances. It is implemented by (*a*) generating high-quality object proposals using YOLOv3, (*b*) computing descriptors of these proposals using an approach based on maximal activation of convolutions,

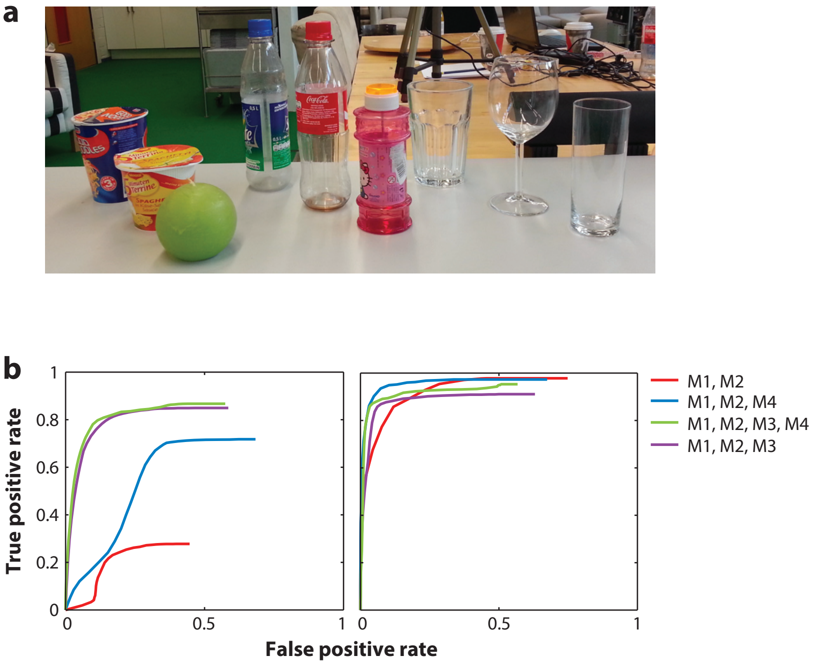

Figure 6

(*a*) A data set of different types of objects. (*b*) Performance graphs for recognizing those objects using four modalities: 2-D shapes (maximum intensity gradients, M1), 3-D geometry (maximum normal vectors, M2), transparency (unavailable depth, M3), and specular reflection (maximum intensity, M4). The left graph shows the receiver operating characteristic (ROC) curve for the entire set of objects, and the right graph shows the curve for a reduced set with no transparent objects. Figure adapted from Reference 78 with permission.

(*c*) recognizing the object instances using an open-set nearest-neighbor classifier, and (*d*) filtering any overlapping recognitions. YoloSPoC outperformed the existing methods in recognizing multiple objects, occlusions, illumination changes, cluttered backgrounds, nontextured objects, and object classes that were unavailable when training the proposal generator. To cope with illumination changes, Houliston et al. (81) proposed a fast method for adapting lookup tables to lighting changes in real time. The method adjusts the classified color space regions while keeping both their surface area and volume constant.

Pepper, one of the standard platforms for the social communication task in the RoboCup@Home domain, has difficulty self-localizing in large environments using its lidars and RGB-D (red, green, and blue plus depth) camera. Gomez et al. (82) proposed a localization and navigation system based on visual simultaneous localization and mapping (SLAM). Furthermore, Schneider et al. (83) proposed gesture recognition in RGB videos using human body keypoints and dynamic time warping for the social communication task.

6.2. Action

The standard platform for the Four-Legged League was the Sony AIBO. Kohl & Stone (84, 85) proposed a machine learning approach for optimizing a quadrupedal trot gait for forward speed. After approximately three hours of learning, the robots achieved a gait faster than any previously known gait for the AIBO, significantly outperforming a variety of existing hand-coded and learned

Supplemental Material >

solutions. This gait was used by many teams (see **Supplemental Figure 12**), and further improvements for subsequent generations of AIBO were achieved through the application of optimization and learning methodologies.

Posturally stable bipedal walking is one of the central issues in the Humanoid League and the current SPL. Sugihara et al. (86) developed stable online walking trajectories using an inverted pendulum model. Many teams in the Humanoid League have developed omnidirectional walking engines and stable walking gaits in forward directions. Hemker et al. (87) developed a sequential surrogate optimization approach that enabled very fast learning and stable forward walking motion with only a few experiments. To develop a stable and adaptive behavior skill, the combination of real robot experiments and simulations is a powerful tool for teams. Rodriguez et al. (88) proposed an approach that combines simulations and real experiments to learn gait stabilization parameters. They used a Bayesian optimization method to select the most informative points in a parameter space to evaluate, based on the entropy of the cost function to optimize. Zahn et al. (89) optimized robot movements, specifically walking and kicking, using genetic algorithms and simulations. For the kick script, the resulting optimal configuration improved the kick distance by a factor of six, with 50% less torso sway. For the walk engine, the forward speed increased by 50%, with 38% less torso sway, as compared with a manually tuned walk engine. **Figure 7** shows the results from two solutions derived from the genetic algorithm.

Because the hardware platform is fixed in the SPL, teams have devised several methods for improving walking and kicking using different software. Iverach-Brereton et al. (90) proposed a method for enabling the humanoid robot to balance on highly dynamic terrains using fuzzy logic on a humanoid DARwIn-OP robot. Böckmann & Laue (91) applied a popular dynamic-movement-primitives approach to the domain of soccer-playing humanoid robots to obtain a kick

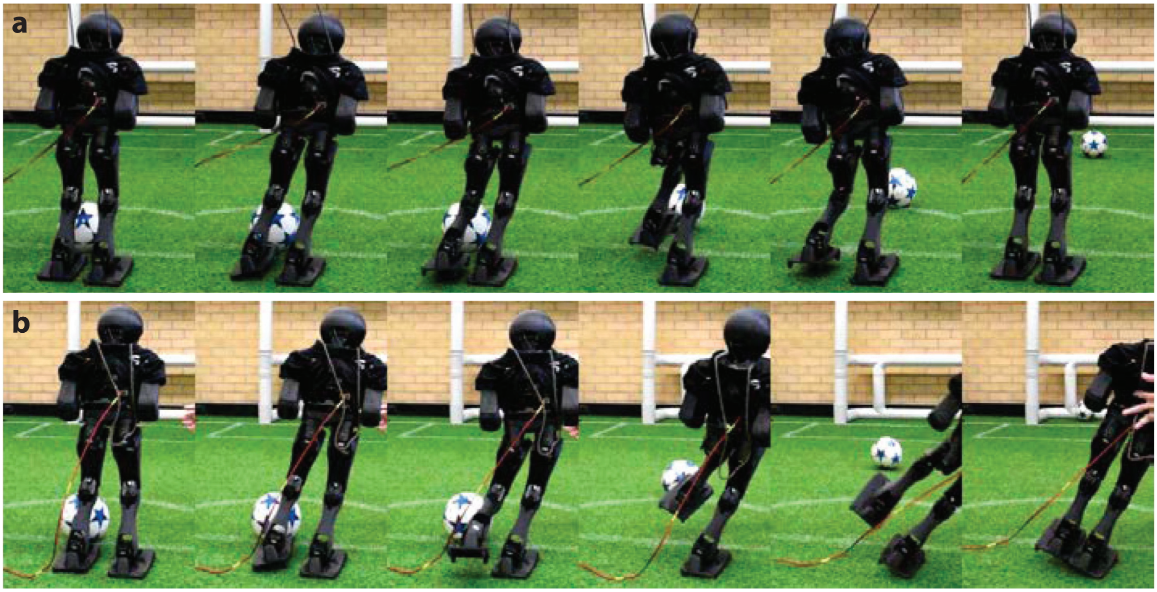

Figure 7

Results from two solutions based on the genetic algorithm of Zahn et al. (89). Panel *a* shows one set of optimized movements; although the distance covered is slightly shorter than that of panel *b*, the stability is slightly better. In real situations, the configuration in panel *b* led to the robot falling over; in the simulations, however, the robot was able to remain standing after either set of movements. Figure adapted from Reference 89.

motion for the NAO robot. This approach includes a mathematical motor model that compensates for the NAO robot's motor-control delay, as well as a novel minor extension to the dynamic-movement-primitives formulation. Seekircher & Visser (92) utilized a linear-inverted-pendulum-based closed-loop walk model that adapts to differences in the physical behavior of the robot by optimizing parameters of the model directly on the NAO while walking and executing other tasks (see **Supplemental Figure 13**).

Supplemental Material >

In the 3-D soccer simulation, the NAO model is used to improve the actions. Masterjohn et al. (93) investigated the decision-making and behavior of robotic biped goalkeepers and proposed two approaches: a heuristics-based approach with linear regression and Kalman filters for improved perception, and another approach based on mental models with nonlinear regression for ball trajectory filtering. Based on various simulations, they concluded that both approaches would significantly improve the goalkeepers' save success rates. Lanari et al. (94) tackled the gait-planning problem by using a flexible linear-inverted-pendulum model. They extended a stable-inversion approach to obtain bounded center-of-mass reference trajectories, and this approach showed several advantages over preview control. Abdolmaleki et al. (95) designed a flexible kick controller that controls the robot (nearly) optimally for a continuous range of kick distances, based on a contextual policy search method. Kasaei et al. (96) achieved a forward velocity of 80.5 cm/s after optimizing the parameters using a genetic algorithm. Peña & Visser (97) proposed a walk–kick framework that can generate a kick trajectory in an arbitrary direction without prior input or knowledge of the parameters of the kick in the midst of walking, while still guaranteeing that a reference trajectory is achieved.

In the RoboCup@Home domain, Mitrevski et al. (98) analyzed dynamic motion primitives in the context of a Toyota HSR and extended the primitives to make it possible to perform a whole-body motion using a mobile manipulator. Renault et al. (99) analyzed the literature on navigation among movable obstacles and found that social acceptability remains an unaddressed problem in this robotics navigation domain. They developed a simulator that allowed testing of their social mobility evaluations for obstacle selection and social placements of objects using a semantic map layer.

Ball interception is a necessary skill for switching from defense mode to offense mode. Makarov et al. (100) proposed a model-free algorithm for intercepting a moving ball using a geometric approach. Two key ideas are the consideration of ball motion via a transition to a reference frame where the ball is static, and planning the motion of the robot in such a reference frame from a geometric viewpoint. The method successfully achieved ball interceptions in a variety of scenarios in the SSL competitions.

For mobile robot control in the RoboCupSoccer SSL and RoboCupRescue Robot League, Ommer et al. (101) proposed a new adaptive compensation feedforward controller. This controller is capable of learning a compensation motion model online without any prior knowledge, so as to counteract nonmodeled disturbances such as slippages or hardware malfunctions.

Omnidirectional motion control is very popular in the SSL, MSL, and other leagues. However, the problem of time-optimal control of omnidirectional robots with bounded acceleration (TOC-ORBA) remains unsolved. Balaban et al. (102) proposed a real-time solver for true TOC-ORBA. They introduced a two-stage optimal control solver and implemented it in a real robot in the SSL to verify the efficiency of the solver.

7. ROBOT COGNITION AND LEARNING

The robot cognition and learning category focuses on the processing of more global and longer-term information, whereas the perception and action category focuses on real-time (immediate) processing. Machine learning, especially deep learning, is a good tool for processes involving reinforcement learning (RL) methods. Traditional AI approaches are also useful.

RL has been used from the inception of RoboCup, and methods have been continuously devised to avoid or address several issues in RL. The two main issues are the curse of dimensionality in the action and state spaces and the reduction of computation time. To cope with these issues, Lobos-Tsunekawa et al. (103) proposed the use of decentralized RL with finite support basis functions as an alternative to a Gaussian radial basis function. As a testbed, they used an RL-based controller for a midwalk kick with NAO robots. Compared with classical approaches, this method saved up to 99.94% of execution time and 98.82% of memory consumption during execution without diminishing performance.

Another issue in RL is the design of the reward function. Watkinson & Camp (104) introduced the use of transfer learning. They demonstrated the possibility of training an agent through a series of increasingly difficult tasks with fewer training iterations rather than using engineered rewards. This approach seems similar to the learning-from-easy-missions approach (105).

Running is a significant challenge, even in a 3-D simulation of the NAO robot. Abreu et al. (106) proposed a way of leveraging a proximal policy optimization using the information provided by the simulator for official RoboCup matches. By using a mix of raw, computed, and internally generated data, they achieved a sprinting speed of approximately 2.5 m/s. Both the sprinting and stopping behaviors were remarkably stable (see **Supplemental Figure 14**).

Supplemental Material >

Collision avoidance is another germane issue for indoor service robots. Leiva et al. (107) proposed an end-to-end approach to endow indoor service robots with the ability to avoid collisions using deep RL. Their approach enabled a robot to learn a proficient collision avoidance policy from scratch (see **Supplemental Figure 15**).

In the RoboCupRescue domain, because teleoperation requires certain skills for maneuverability on irregular surfaces, some sort of autonomy is desirable. Wiley et al. (108) developed a system that learns the effects of a robot's actions and then uses this knowledge to plan an approach to reconfiguring the robot's tracks so that it can overcome different types of obstacles. The system is a hybrid of qualitative symbolic learning and RL. High-level knowledge regarding the task could reduce the number of attempts necessary to learn a new skill.

A number of general learning approaches have been proposed. Rizzi et al. (109) proposed a situation-aware fear learning (SAFEL) model and discussed specific scenarios where SAFEL's associative learning could help to increase the positive outcomes of a team during a soccer match through contextual adaptation. SAFEL enables NAO robots in the SPL to learn the behavioral profile of the opposing team at runtime. Simoes et al. (110) observed that the weighted-policy-learner algorithm has difficulty regarding convergence to deterministic strategies, and they proposed an adjusted and bounded weighted policy learner with a new update rule. In this new rule, the algorithm's speed is not slowed, and its behavior in stochastic Nash equilibrium games remains unchanged.

In the robot task-learning methods, the most typical representation of a state and action is a state machine representation, in which states encapsulate actions and a transition function switches between states. In real robot experiments or competitions, roboticists often adjust the parameters manually to cope with any changes in the environment. Holtz et al. (111) proposed a semiautomatic white-box approach for adjusting the transition parameters of robot state machines. The proposed method effectively increased the success rate for multiple behaviors, such as finding new parameters quickly using a small number of annotations, producing solutions that generalize well to novel situations, and improving the performance in a real-world robot soccer application, the RoboCup SSL.

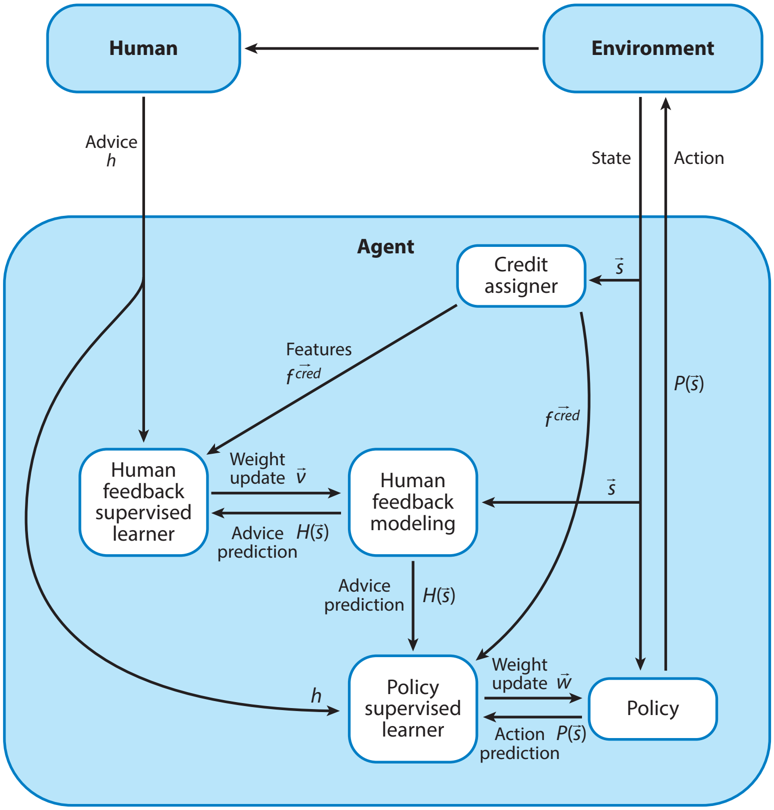

Figure 8

Block diagram of the Corrective Advice Communicated by Humans (COACH) learning framework. Figure adapted from Reference 112 with permission.

8. HUMAN–ROBOT INTERACTION

Apart from the soccer competitions, daily-life (RoboCup@Home) and special (RoboCupRescue) applications have become increasingly important. In these situations, human–robot interaction provides rich and deep research issues.

Celemin & Ruiz-del-Solar (112) proposed an interactive learning framework called Corrective Advice Communicated by Humans (COACH) that allows nonexperts to shape a policy through corrective advice, which includes a mechanism for adaptively adjusting the amount of human feedback that a given action receives, considering past feedback. **Figure 8** shows a block diagram of the COACH learning framework. Celemin & Ruiz-del-Solar (112) found that COACH outperformed existing frameworks.

Facial analysis techniques have become a crucial component of human–machine interaction in the fields of assistive and humanoid robotics. However, robustness against variations in head pose is a substantial challenge. Grupp et al. (113) proposed a real-time-capable 3-D face-modeling framework for 2-D in-the-wild images using a fully automatic landmark-based approach for fitting a 3-D morphable model. They performed real-time processing using in-the-wild images that serve as a preprocessing method for various facial analysis tasks.

Domestic robots must be able to collect various information regarding people to perform tasks and conduct socially acceptable human–robot interactions. Saraydaryan et al. (114) proposed a framework for extracting high-level human features from a 2-D camera in addition to tracking people over time. Recognition of people's poses and postures, clothing colors, and faces is combined with tracking and reidentification abilities. This framework was successfully used with a Pepper robot in the 2018 RoboCup@Home competition.

The 2018 champion team in the RoboCup@Home Domestic Standard Platform League, Hibikino-Musashi@Home, developed a very-large-scale integration (VLSI) chip based on the Time-Domain Analog Computing with Transient States (TACT) approach for intelligent processing on robots. This chip was integrated into a robot via ROS interfaces. The team demonstrated a human-tracking robot and received the Best Live Demonstration Award at the 2019 IEEE International Symposium on Circuits and Systems (115). They also implemented a brain-inspired amygdala model in hardware and applied the proposed amygdala model to a robot waiter task in a restaurant. In this experiment, the model learned a customer's preferences after only a few human–robot interactions, outperformimg a software implementation on an Intel Core i5-3470 CPU (116).

Linguistic communication is essential for domestic robots to assist humans at home. Specifically, it is necessary for robots to connect natural language to the physical world to manipulate objects through a reasoning process. Lu & Chen (117) proposed an architecture that combines grounding and planning to enable robots to solve such a problem. The grounding system depends on the robot's sensors and generation of a knowledge base for the physical environment. The planning system utilizes the knowledge base to infer a plan for object manipulation.

In the case of human–robot language communication, the robot often cannot understand human commands because of the robot's uncertainty of the situation and human misunderstanding of the robot's comprehension of the situation. Gemignani et al. (118) addressed the problem of allowing a human to understand a robot's internal representation through dialogue. They introduced the concept of sensing descriptors, which the robot uses to recognize unknown properties in the human's commands and then warn the human about them. Unknown properties can be learned over time by leveraging past interactions to enhance the grounding capabilities of the robot (see **Supplemental Figure 16**).

Supplemental Material >

Matamoros et al. (119) addressed the language communication performance levels of the RoboCup@Home domain in general and created a pipelined road map for stimulating research in the area of natural language understanding as it applies to domestic service robotics. Semantic parsing is one way of converting natural language commands into executable representations. However, the current semantic parsing has an application limitation. To address this issue, Walker et al. (120) proposed an approach that leverages neural semantic parsing methods in combination with contextual word embedding to enable the training of a semantic parser with few data and without domain-specific parser engineering. Their results show that neural semantic parsers can predict the logical form of unseen commands with 89% accuracy.

Regarding the issue of architecture in the RoboCup@Home domain, Jumel et al. (121) proposed an architecture dedicated to the orchestration of high-level abilities for humanoid robots such as Pepper. The architecture was required to perform tasks similar to the ones proposed in the RoboCup@Home competitions. Context awareness is a key feature of their system. Peña et al. (122) discussed a modular agent architecture for an interactive system that integrates two frameworks (an in-house virtual social agent and a robot agent framework) and enables social multimodal human–robot interaction with the Toyota HSR. Their pilot study revealed no significant differences in enjoyment, friendliness, competence, uncanniness, and other categories when comparing Toyota HSRs with and without an embodied empathetic virtual agent (eEVA). They concluded that the eEVA's character does not make the Toyota HSR more uncanny, boring, or annoying.

9. MULTIROBOT SYSTEMS

Multiagent teamwork is one of the most essential issues in RoboCup. How can collaborative and competitive behaviors be explicitly preprogrammed or made to implicitly emerge through learning? To decide whether to pass a ball to a teammate (expecting it to either shoot or make another pass) or to shoot the ball itself, a robot needs to estimate the capabilities of its teammates in a dynamically changing environment. The 2-D Simulation League, SSL, and MSL are considering this issue, and the order here indicates the increased difficulty in addressing it.

In the 2-D Simulation League, the evaluation functions used for the decision-making process are among the most influential factors. Fukushima et al. (123) proposed a method that improves the performance of a team by mimicking a stronger team. They employed a neural network to model an expert team's evaluation function using positive and negative episodes of action sequences that are extracted from game logs, and the proposed method successfully improved the performance (e.g., win rate and scored goals) of their team. Wiretapping has been permitted in this league from the beginning, although its main purpose is to share information among teammates. Gabel et al. (124) proposed an approach to wiretapping and decoding opponent communication and systematically evaluating its impact. Consequently, a team that wiretaps its opponent and exploits intercepted information appropriately can significantly boost its own playing performance.

Ball manipulation for an individual robot already in possession of the ball is a typical issue in the SSL. The robot must intelligently move the ball to its target destination while keeping it away from opponents. Cooksey et al. (125) presented and compared complementary ball manipulation skills and described an approach to selecting the appropriate skill given the situation. Adachi et al. (126) proposed a method for identifying strategies by classifying an observed sequence of basic actions selected by an opponent during a game. The strategies of their own team are evaluated against those of four opponent teams using a Rand index, and the resultant value of 0.877 (where >0.840 is a high value) indicates a high level of classification algorithm performance. A skills, tactics, and plays (STP) architecture was developed by the CMDragons team and was very popular for many years. Schwab et al. (127) demonstrated how modern deep RL techniques can be incorporated into an existing STP architecture. They used the deep deterministic policy gradient algorithm to learn skills, compared the learned skills with existing ones, and demonstrated how RL can be leveraged to learn simple skills that can be fused by humans into high-level tactics that allow an agent to navigate a ball, aim, and shoot a goal. They also found that the positioning of the opponent's team becomes increasingly important as the SSL game increases in complexity. Laureano & Tonidandel (128) proposed the use of a particle swarm optimization algorithm as an option for determining the positioning during a match and demonstrated the feasibility of applying this algorithm to finding the robots' positions.

In an adversarial multiagent environment, a balance between the advantages and disadvantages of completely decentralized solutions and centralized ones is a key issue, and RoboCupSoccer typically provides such situations. Dias et al. (129) weighed in on this issue and proposed the solution of electing a leader from among the robots on the team. Their proposed solution builds on the Raft algorithm, which has two limitations: It fails to elect a leader when fewer than three nodes are available, and it does not prioritize among candidates. To overcome these limitations, they adopted a backup system and a preferred leader agent (see **Supplemental Figure 17**). Their team, CAMBADA (Cooperative Autonomous Mobile Robots with Advanced Distributed Architecture), applied this method in their team play (130). More centralized architectures were used by the Tech United Eindhoven (131) and Water (132) teams, both of which showed brilliantly performed passes at the MSL final of RoboCup 2019. Another issue in an adversarial multiagent environment is circumnavigation control, such as entrapping a hostile target. Yao et al.

Supplemental Material >

(133) proposed distributed circumnavigation control with dynamic spacing for a heterogeneous multirobot system. They introduced utility and formation guidelines to address dynamic spacing according to the robots' properties and presented a theoretical analysis using graph theory along with experiments to prove the effectiveness of the proposed algorithm based on utilities.

10. EDUCATION AND EDUTAINMENT

RoboCup is a proven excellent model for conducting project-based learning that provides young students with experiences and training in fundamental STEM subjects in very practical situations and helps them understand how robots work through the competitions. The RoboCupJunior domain focuses on creative robot performance (e.g., dancing) and evaluates the total performance from a variety of perspectives, such as artistic impressions for appearance and movements.

Wong et al. (134) reported the success of their activities in the Hunter region of New South Wales, Australia, in 2012 and 2017 in terms of the number of participants and their fields, STEM scores, and gender balance. Moreover, they reported the high potential of the older boys and girls to proceed to future activities in science and technology. Hughes et al. (135) reviewed robot rescue simulation platforms for robotics education, focusing on a natural learning curve to provide appropriate rescue challenges for different age groups. They discussed the requirements for such a platform and compared several different platforms. They concluded that the case study of a sample game-field rescue simulation platform was suitable for students at different points along the learning curve.

Recently, several new activities have strengthened the collaboration between the major and junior leagues: the Rapidly Manufactured Robot Challenge (a RoboCupJunior Rescue competition) and new entry levels in the RoboCupRescue Simulation League, RoboCup@Home domain, and RoboCupSoccer Humanoid League.

11. APPLICATIONS AND BENCHMARKING

Another unique aspect of RoboCup is in the design of performance evaluation. The game itself is a good indicator of the total performance on the system level. Nevertheless, each league has its own technical challenges to evaluate achievements on module and functionality levels toward the league's goal. These benchmarks are good indicators of whether the technologies fostered in RoboCup are broadly applicable across different fields.

The robot hardware and software constitute the infrastructure for any robotics study. In the case of RoboCup, research on robot hardware and software appears to be geared specifically toward a final goal. However, this work does not merely facilitate the application of technologies fostered in RoboCup to tasks aimed at achieving RoboCup's ultimate goal. Even partially achieved technologies, such as those of Kiva Systems (now Amazon Robotics), can potentially be channeled to practical applications at any point in their development.

11.1. Challenge and Task Design

Both of the RoboCupIndustrial leagues (RoboCup@Work and the RCLL) focus on industrial application (automated reasoning and planning, and mobile manipulation), with slightly different specifications. Zug et al. (136) proposed a crossover challenge to foster closer cooperation between the two leagues. They outlined four integration milestones and proposed a specific scenario and task for the first milestone.

Making progress toward the ultimate goal requires more researchers and universities that might bring advanced technologies and new perspectives to research. And attracting new researchers requires an easy entry path for new teams that is suitable for undergraduate students and universities with limited budgets. Gerndt et al. (137) proposed an entry-level league with a reduced set of requirements that can bridge the gap between the junior level and the advanced robots and teams in the Humanoid League. They called it the "Humanoid Rookie (Sub-) League (HRL)" and suggested that it would give new researchers and teams the time to gather the experience and funds needed to successfully participate in and contribute to the Humanoid League's progress toward the 2050 goal. This proposal addresses the general challenge of bridging the gap between the highly advanced teams and robots in the major leagues and the resulting high and continuously rising entry level for newer teams.

Generally, robots require long-term memory for coherent interaction and communication with humans. In the RoboCup@Home domain, a robot should be a continuous being in order to maintain long-term relationships with humans. Pavez et al. (138) observed that such a capability is essential for a home robot to perform some domestic tasks and proposed a new challenge: the verification of the capability of long-term memory in robots.

After the Fukushima Daiichi nuclear power plant disaster, the RoboCupRescue Robot League changed the regulations of its competitions to reflect real disaster situations as much as possible. This trend became mainstream in the league, and firefighters, police personnel, and other responders are now invited from all over the world. At RoboCup 2019 in Sydney, Australia, the New South Wales Police Rescue and Bomb Disposal Unit played a vital role in preparing the arena. They also brought their brand-new robots to test them within the arena and exhibit them to the researchers and students (see **Supplemental Figure 18**). Observing these changes, Shimizu & Takahashi (139) proposed a new standard task based on ordinary tasks in the RoboCupRescue Simulation League. They surveyed the rescue competitions and realized the potential of the virtual league to imitate actual situations.

Supplemental Material >

11.2. Performance Check

It is necessary to evaluate RoboCup games to tackle all the germane issues, such as the multiagent team strategy. A 2-D simulation league is suitable for this purpose owing to the rich game data available. Gabel et al. (140) evaluated the game quality across the past 20 years, observing that although the first decade showed amazing progress, the second did not. Michael et al. (141) proposed a new approach for identifying situations and behaviors. Their goal was to identify situations from data in an unsupervised way without making use of predefined soccer-specific concepts, such as passing or dribbling. The system can segment games into sequences of situations that are learned in an unsupervised way and learn conceptors that are useful for predicting the near future of each situation. Suzuki & Nakashima (142) proposed a forward simulation for situation evaluation (FOSSE) approach for evaluating game situations. FOSSE generates future game situations using forward simulation. Computational experiments were conducted to verify the effectiveness of the proposed approach.

Because game evaluation can be enhanced by using richer data with less noise, data generation is a powerful tool for that purpose. Michael et al. (143) generated game data with incomplete and noisy percepts (as sent to each player) in addition to a ground-truth log file created by the simulator (global, complete, noise-free information on all objects on the field). These data were made available as comma-separated value (CSV) files as well as in the original soccer simulator formats. Pomas & Nakashima (144) proposed a CNN that assesses a situation at one point of a RoboCup 2-D soccer game and predicts which team will score next and when using only the soccer field images as input. The next goal is predicted using SituationScores, which estimates the remaining number of frames; the average error of SituationScores was less than that of the existing methods. SituationScores are also used for soccer monitoring, to make the experience of watching games more entertaining (145). Fukushima et al. (146) discussed the validity of similarity measures for action trajectories based on kick distributions, which were used to estimate the dissimilarity (or distance) between the strategies of two teams. They demonstrated that the similarity analysis methods have a positive correlation with human subjectivity, implying that their method is valid for the similarity analysis. In addition, the calculation time could be reduced by using continuous kick probability distributions (see **Supplemental Figure 19**).

Supplemental Material >

Ball detection is widely studied in relation to RoboCup, as mentioned in Section 6.1, as well as in other studies. Gabel et al. (147) proposed a patch-based approach and evaluated the ball detection performance at RoboCup 2017 using several methods, such as Harr, AlexNet, and Inception, with different steps and epochs. These networks are for generic use and can thus be modified for RoboCup specifications. Speck et al. (148) proposed a state-of-the-art ball detection model and made training and test data sets (comprising more than 35,000 and more than 2,000 images, respectively) publicly available for teams to train the models under real-time requirements and with the limited computing power of humanoid robots. They frequently release new versions to enable teams to hone their performance.

In the RoboCup@Home domain, the number of objects that the robot needs to recognize and manipulate is much larger than that of the RoboCupSoccer domain. Recent progress in deep learning, especially in CNNs, has enabled the recognition of various types of objects. However, the data set for RoboCup@Home use is limited. Massouh et al. (149) proposed a benchmark for object recognition, including a large-scale training set of 196,000 images labeled with classes derived from RoboCup@Home rule books (the RoboCup@Home-Objects data), two medium-scale test sets (one taken with a Pepper robot) with different objects and different backgrounds with respect to the training set, a robot behavior for image acquisition, and several useful analyses of the results. The RoboCup@Home-Objects data are very useful, not only for the teams participating in RoboCupSoccer but also for the technical committee designing and evaluating the competition.

Natural language communication is another issue in the RoboCup@Home general-purpose service robot category, which involves supporting humans in the context of domestic tasks. Kramer et al. (150) analyzed a comparative study and proposed the benchmarking of four natural language understanding models, called Mbot, Rasa, LU4R, and ECG. They evaluated the four models in the competitions and concluded that both Mbot and Rasa are suitable for robot command understanding; however, Mbot is slightly more suitable, as Rasa has difficulty differentiating between certain location entities, such as destination and source categories.

The RCLL generates a massive amount of data regarding the state changes of the game and communication with the robots. Niemueller et al. (151) analyzed the data from the 2014 competition through key performance indicators. Applying adapted key performance indicators to the RCLL provides interesting insights regarding the strategies of the robot teams. In progressing toward more realistic industrial properties with 24/7 production, teams should perform in shifts (i.e., without an intermediate environment reset).

12. CONCLUDING REMARKS

Now in its 24th year, RoboCup is still unique in multiple ways. It was the first organization to introduce competitions and benchmarking through formidable and publicly appealing visionary challenges for research in robotics and AI. The initial vision of a team of humanoid robots that are physically and mentally on par with a World Cup champion human soccer team has been extended to the deployment of intelligent robots that can directly address societal challenges and disaster response (with Industry 4.0 and Society 5.0 also being forerunners in these domains). RoboCup is supported and self-funded by a large and truly international community consisting of approximately 1,000 and 10,000 major and junior teams, respectively, that work cooperatively and enthusiastically toward the common goals of RoboCup. RoboCup successfully eliminates the boundaries between fundamental research and evaluation and benchmarking on a full-system level, between organizers and participants, between academia and industry, between university-level research and project-oriented STEM education, between mere short-term acceleration of research and development through robotics competitions and their long-term sustainability, and—last but not least—between researchers, students, and teachers with different cultural backgrounds from approximately 45 nations and regions around the world.

This article has provided an overview of research activities from the perspective of RoboCup, focusing mainly on selected studies from recent years. Toward the realization of RoboCup's ultimate goal, the unique cycle of deliberation (new idea or challenge), implementation (design and realization), and verification (competition) will continue to involve new teams, researchers, and supporters. In the annual RoboCup event, which is open to the public, the large variety of test fields and competitions developed can be observed, along with the high level of enthusiasm of researchers and students in the competitions. The most recent RoboCup was held in Sydney in 2019, and the next two will take place in 2020 in Bordeaux, France, and in 2021 in Bangkok, Thailand.

Although RoboCup already incorporates an amazing variety of research issues from different disciplines of robotics and AI, it could further benefit from stronger incorporation of additional topics. These topics include (*a*) higher-level cognitive architectures with features such as self-consciousness, attention, and emotion (e.g., 152, 153) to help robots achieve their mission; (*b*) compliant robotics technology (soft robotics) to enable more dynamic motions, such as running and jumping, with human-like dexterity, in addition to greater energy efficiency in humanoid robots; and (*c*) the use of neuromorphic chips and devices for energy conservation and efficient computation [e.g., the Neuromorphic Dynamics Project (**http://www.ams.eng.osaka-u.ac.jp/nedo-nmd/?lang=en**)], which only a few teams have attempted to implement (115, 116).

DISCLOSURE STATEMENT

Both authors are members of the board of trustees of the RoboCup Federation.

ACKNOWLEDGMENTS

The authors express their unreserved appreciation for the inspiration provided by helpful discussions with the members of the board of trustees, the executive committee, and the technical committees of the leagues of the RoboCup Federation in the direct preparation of this article and beyond. They also thank Jim Duncan and the Annual Reviews team for their support in publishing this article.

LITERATURE CITED

1. Kitano H, Asada M, Kuniyoshi Y, Noda I, Osawa E. 1995. RoboCup: the Robot World Cup Initiative. In *Proceedings of IJCAI-95 Workshop on Entertainment and AI/Alife*. **https://www2.sonycsl.co.jp/person/kitano/RoboCup/RoboCup.ps**
2. Kitano H, Asada M, Kuniyoshi Y, Noda I, Osawa E, Matsubara H. 1998. RoboCup: a challenge problem for AI and robotics. See Ref. 10, pp. 1–19
3. Asada M, Kitano H, Noda I, Veloso M. 1999. RoboCup: today and tomorrow—what we have learned. *Artif. Intell.* 110:193–214

4. Stone P, Quinlan M, Hester T. 2010. The essence of soccer, can robots play too? In *Soccer and Philosophy: Beautiful Thoughts on the Beautiful Game*, ed. T Richards, pp. 75–88. Chicago: Open Court
5. Nagatani K, Kiribayashi S, Okada Y, Otake K, Yoshida K, et al. 2013. Emergency response to the nuclear accident at the Fukushima Daiichi Nuclear Power Plants using mobile rescue robots. *J. Field Robot.* 30:44–63
6. Nardi D, Noda I, Ribeiro F, Stone P, von Stryk O, Veloso M. 2014. RoboCup soccer leagues. *AI Mag.* 35(3):77–85
7. RoboCup. 2019. SSL-Vision wiki. *GitHub*. **https://github.com/RoboCup-SSL/ssl-vision/wiki**
8. Gerndt R, Seifert D, Baltes J, Sadeghnejad S, Behnke S. 2015. Humanoid robots in soccer: robots versus humans in RoboCup 2050. *IEEE Robot. Autom. Mag.* 22(3):147–54
9. Noda I, Suzuki S, Matsubara H, Asada M, Kitano H. 1998. Overview of RoboCup-97. See Ref. 10, pp. 20–41
10. Kitano H, ed. 1998. *RoboCup-97: Robot Soccer World Cup I*. Berlin: Springer
11. Asada M, Kitano H, eds. 1999. *RoboCup-98: Robot Soccer World Cup II*. Berlin: Springer
12. Veloso M, Pagello E, Kitano H, eds. 2000. *RoboCup-99: Robot Soccer World Cup III*. Berlin: Springer
13. Stone P, Balch T, Kraetzschmar G, eds. 2001. *RoboCup 2000: Robot Soccer World Cup IV*. Berlin: Springer
14. Birk A, Coradeschi S, Tadokoro S, eds. 2002. *RoboCup 2001: Robot Soccer World Cup V*. Berlin: Springer
15. Kaminka G, Lima PU, Rojas R, eds. 2003. *RoboCup 2002: Robot Soccer World Cup VI*. Berlin: Springer
16. Bonarini A, Polani D, Browning B, Yoshida K, eds. 2004. *RoboCup 2003: Robot Soccer World Cup VII*. Berlin: Springer
17. Nardi D, Riedmiller M, Sammut C, Santos-Victor J, eds. 2005. *RoboCup 2004: Robot Soccer World Cup VIII*. Berlin: Springer
18. Bredenfeld A, Jacoff A, Noda I, Takahashi Y, eds. 2006. *RoboCup 2005: Robot Soccer World Cup IX*. Berlin: Springer
19. Lakemeyer G, Sklar E, Sorrenti DG, Takahashi T, eds. 2007. *RoboCup 2006: Robot Soccer World Cup X*. Berlin: Springer
20. Visser U, Ribeiro F, Ohashi T, Dellaert F, eds. 2008. *RoboCup 2007: Robot Soccer World Cup XI*. Berlin: Springer
21. Iocchi L, Matsubara H, Weitzenfeld A, Zhou C, eds. 2009. *RoboCup 2008: Robot Soccer World Cup XII*. Berlin: Springer
22. Baltes J, Lagoudakis MG, Naruse T, Ghidary SS, eds. 2010. *RoboCup 2009: Robot Soccer World Cup XIII*. Berlin: Springer
23. Ruiz-del-Solar J, Chown E, Ploger PG, eds. 2011. *RoboCup 2010: Robot Soccer World Cup XIV*. Berlin: Springer
24. Roefer T, Mayer N, Savage J, Saranl U, eds. 2012. *RoboCup 2011: Robot Soccer World Cup XV*. Berlin: Springer
25. Chen X, Stone P, Sucar L, van der Zant T, eds. 2013. *RoboCup 2012: Robot Soccer World Cup XVI*. Berlin: Springer
26. Behnke S, Veloso M, Visser A, Xiong R, eds. 2014. *RoboCup 2013: Robot Soccer World Cup XVII*. Berlin: Springer
27. Bianchi R, Akin H, Ramamoorthy S, Sugiura K, eds. 2015. *RoboCup 2014: Robot Soccer World Cup XVIII*. Cham, Switz.: Springer
28. Almeida L, Ji J, Steinbauer G, Luke S, eds. 2016. *RoboCup 2015: Robot Soccer World Cup XIX*. Cham, Switz.: Springer
29. Behnke S, Sheh R, Sarıel S, Lee DD, eds. 2017. *RoboCup 2016: Robot Soccer World Cup XX*. Cham, Switz.: Springer
30. Obst HAO, Sammut C, Tonidandel F, eds. 2018. *RoboCup 2017: Robot Soccer World Cup XXI*. Cham, Switz.: Springer
31. Holz D, Genter K, Saad M, von Stryk O, eds. 2019. *RoboCup 2018: Robot Soccer World Cup XXII*. Cham, Switz.: Springer
32. Suthakorn J, Williams M, Niemueller T, Chalu S, eds. 2020. *RoboCup 2019: Robot Soccer World Cup XXIII*. Cham, Switz.: Springer

33. Douven Y, Houtman W, Schoenmakers F, Koen Meessen HVDL, Bruijnen D, et al. 2019. Tech United Eindhoven Middle Size League winner 2018. See Ref. 31, pp. 413–24
34. Junkai R, Chenggang X, Junhao X, Kaihong H, Huimin L. 2016. A control system for active ball handling in the RoboCup Middle Size League. In *2016 Chinese Control and Decision Conference*, pp. 4396–402. Piscataway, NJ: IEEE
35. Beuermann M, Ossenkopf M, Geihs K. 2020. Positioning of active wheels for optimal ball handling: a guide for designing a new ball handle mechanism for Middle-Size-League at RoboCup. See Ref. 32, pp. 30–43
36. Gies V, Soriano T, Albert C, Prouteau N. 2020. Modelling and optimisation of a RoboCup MSL coilgun. See Ref. 32, pp. 71–85
37. Yoshimoto T, Horii T, Mizutani S, Iwauchi Y, Yamada Y, et al. 2017. *OP-AmP 2017 team discription paper*. Descr. Pap., Asagami Works, Osaka, Jpn. **https://ssl.robocup.org/wp-content/uploads/2019/01/2017_TDP_Op-Amp.pdf**
38. Yoshimoto T, Horii T, Mizutani S, Iwauchi Y, Zenji S. 2019. *OP-AmP 2019 extended team discription paper*. Descr. Pap., Asagami Works, Osaka, Jpn. **https://ssl.robocup.org/wp-content/uploads/2019/03/2019_ETDP_OP-AmP.pdf**
39. Ha I, Tamura Y, Asama H, Han J, Hong DW. 2011. Development of open humanoid platform DARwIn-OP. In *SICE Annual Conference 2011*, pp. 2178–81. Piscataway, NJ: IEEE
40. Fabre R, Rouxel Q, Passault G, N'Guyen S, Ly O. 2017. Dynaban, an open-source alternative firmware for Dynamixel servo-motors. See Ref. 29, pp. 169–77
41. Bestmann M, Guldenstein J, Zhang J. 2020. High-frequency multi bus servo and sensor communication using the Dynamixel protocol. See Ref. 32, pp. 16–29
42. Schwarz M, Pastrana J, Allgeuer P, Schreiber M, Schueller S, et al. 2014. Humanoid TeenSize open platform NimbRo-OP. See Ref. 26, pp. 568–75
43. Farazi H, Ficht G, Allgeuer P, Pavlichenko D, Rodriguez D, et al. 2019. NimbRo robots winning RoboCup 2018 Humanoid AdultSize soccer competitions. See Ref. 31, pp. 436–49
44. Yamamoto T, Takagi Y, Ochiai A, Iwamoto K, Itozawa Y, et al. 2020. Human support robot as research platform of domestic mobile manipulator. See Ref. 32, pp. 457–65
45. Scheunemann MM, van Dijk SG. 2020. ROS 2 for RoboCup. See Ref. 32, pp. 429–38
46. Thielke F, Hasselbring A. 2020. A JIT compiler for neural network inference. See Ref. 32, pp. 448–56
47. Mitrevski A, Ploger PG. 2020. Reusable specification of state machines for rapid robot functionality prototyping. See Ref. 32, pp. 408–17
48. Mellmann H, Schlotter B, Strobel P. 2020. Toward data driven development in RoboCup. See Ref. 32, pp. 176–88
49. Berlin United Nao Team Humboldt. 2019. Tools for data driven research and development in RoboCup. *Humboldt University Berlin*. **https://www.naoteamhumboldt.de/en/projects/RoboCup-data-collection-and-evaluation**
50. Fiedler N, Bestmann M, Hendrich N. 2019. ImageTagger: an open source online platform for collaborative image labeling. See Ref. 31, pp. 162–69
51. Hess T, Mundt M, Weis T, Ramesh V. 2018. Large-scale stochastic scene generation and semantic annotation for deep convolutional neural network training in the RoboCup SPL. See Ref. 30, pp. 33–44
52. Visser A, Nardin LG, Castro S. 2019. Integrating the latest artificial intelligence algorithms into the RoboCup rescue simulation framework. See Ref. 31, pp. 476–87
53. van Dijk SG, Scheunemann MM. 2019. Deep learning for semantic segmentation on minimal hardware. See Ref. 31, pp. 349–61
54. Gholami A, Moradi M, Majidi M. 2020. A simulation platform design and kinematics analysis of MRL-HSL humanoid robot. See Ref. 32, pp. 387–96
55. Inamura T, Mizuchi Y. 2018. Competition design to evaluate cognitive functions in human-robot interaction based on immersive VR. See Ref. 30, pp. 84–94
56. Takami S, Takayanagi K, Jaishy S, Ito N, Iwata K, et al. 2018. Proposed environment to support development and experiment in RoboCup rescue simulation. See Ref. 30, pp. 71–83

57. MacAlpine P, Stone P. 2017. UT Austin Villa RoboCup 3D simulation base code release. See Ref. 29, pp. 135–43
58. Prokopenko M, Wang P. 2020. Gliders2d: source code base for RoboCup 2D soccer simulation league. See Ref. 32, pp. 418–28
59. Fiedler N, Brandt H, Gutsche J, Vahl F, Hagge J, Bestmann M. 2020. An open source vision pipeline approach for RoboCup humanoid soccer. See Ref. 32, pp. 387–96
60. Kohlbrecher S, Meyer J, von Stryk O, Klingauf U. 2011. A flexible and scalable SLAM system with full 3D motion estimation. In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics*, 155–60. Piscataway, NJ: IEEE
61. Kohlbrecher S, Petersen K, Steinbauer G, Maurer J, Lepej P, et al. 2012. Community-driven development of standard software modules for search and rescue robots. In *2010 IEEE International Symposium on Safety, Security, and Rescue Robotics*. Piscataway, NJ: IEEE. **https://doi.org/10.1109/SSRR.2012.6523917**
62. Google Code. 2019. ssl-autonomous-refbox. *Google Code*. **https://code.google.com/p/ssl-autonomous-refbox**
63. Zhu D, Biswas J, Veloso M. 2015. AutoRef: towards real-robot soccer complete automated refereeing. See Ref. 27, pp. 419–30
64. Schoenmakers F, Koudijs G, Martinez CL, Briegel M, van Wesel H, et al. 2014. *Tech United Eindhoven team description 2013: Middle Size League*. Descr. Pap., Eindhoven Univ. Technol., Eindhoven, Neth. **https://www.techunited.nl/media/files/TDP2013.pdf**
65. RoboCup-MSL. 2019. RefBox2015. *GitHub*. **https://github.com/RoboCup-MSL**
66. Huang Z, Chen L, Li J, Wang Y, Chen Z, et al. 2019. RoboCup SSL 2018 champion team paper. See Ref. 31, pp. 401–12
67. Speck D, Barros P, Weber C, Wermter S. 2017. Ball localization for RoboCup soccer using convolutional neural networks. See Ref. 29, pp. 19–30
68. Menashe J, Kelle J, Genter K, Hanna J, Liebman E, et al. 2018. Fast and precise black and white ball detection for RoboCup soccer. See Ref. 30, pp. 45–58
69. Leiva F, Cruz N, Bugueño I, Ruiz-del-Solar J. 2019. Playing soccer without colors in the SPL: a convolutional neural network approach. See Ref. 31, pp. 122–34
70. Teimouri M, Delavaran MH, Rezaei M. 2020. A real-time ball detection approach using convolutional neural networks. See Ref. 32, pp. 323–36
71. Kukleva A, Khan MA, Farazi H, Behnke S. 2020. Utilizing temporal information in deep convolutional network for efficient soccer ball detection and tracking. See Ref. 32, pp. 112–25
72. Felbinger GC, Göttsch P, Loth P, Peters L, Wege F. 2019. Designing convolutional neural networks using a genetic approach for ball detection. See Ref. 31, pp. 150–61
73. Cruz N, Lobos-Tsunekawa K, Ruiz-del-Solar J. 2018. Using convolutional neural networks in robots with limited computational resources: detecting NAO robots while playing soccer. See Ref. 30, pp. 19–30
74. Houliston T, Chalup SK. 2019. Visual mesh: real-time object detection using constant sample density. See Ref. 31, pp. 45–56
75. Szemenyei M, Estivill-Castro V. 2019. Real-time scene understanding using deep neural networks for RoboCup SPL. See Ref. 31, pp. 96–108
76. Szemenyei M, Estivill-Castro V. 2020. ROBO: robust, fully neural object detection for robot soccer. See Ref. 32, pp. 309–22
77. Poppinga B, Laue T. 2020. JET-Net: real-time object detection for mobile robots. See Ref. 32, pp. 227–40
78. Hagg A, Hegger F, Plöger PG. 2017. On recognizing transparent objects in domestic environments using fusion of multiple sensor modalities. See Ref. 29, pp. 3–15
79. Reyes E, Gomez C, Norambuena E, Ruiz-del-Solar J. 2019. Near real-time object recognition for Pepper based on deep neural networks running on a backpack. See Ref. 31, pp. 287–98
80. Loncomilla P, Ruiz-del-Solar J. 2020. YoloSPoC: recognition of multiple object instances by using Yolo-based proposals and deep SPoC-based descriptors. See Ref. 32, pp. 154–65

81. Houliston T, Metcalfe M, Chalup SK. 2016. A fast method for adapting lookup tables applied to changes in lighting colour. See Ref. 28, pp. 190–201
82. Gomez C, Mattamala M, Resink T, Ruiz-del-Solar J. 2019. Visual SLAM-based localization and navigation for service robots: the Pepper case. See Ref. 31, pp. 32–44
83. Schneider P, Memmesheimer R, Kramer I, Paulus D. 2020. Gesture recognition in RGB videos using human body keypoints and dynamic time warping. See Ref. 32, pp. 281–93
84. Kohl N, Stone P. 2004. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *IEEE International Conference on Robotics and Automation, 2004*, Vol. 3, pp. 2619–24. Piscataway, NJ: IEEE
85. Kohl N, Stone P. 2004. Machine learning for fast quadrupedal locomotion. In *Proceedings of the 19th National Conference on Artificial Intelligence*, pp. 611–16. Palo Alto, CA: AAAI Press
86. Sugihara T, Nakamura Y, Inoue H. 2002. Real-time humanoid motion generation through ZMP manipulation based on inverted pendulum control. In *Proceedings: 2002 IEEE International Conference on Robotics and Automation*, Vol. 2, pp. 1404–9. Piscataway, NJ: IEEE
87. Hemker T, Sakamoto H, Stelzer M, von Stryk O. 2009. Efficient walking speed optimization of a humanoid robot. *Int. J. Robot. Res.* 28:303–14
88. Rodriguez D, Brandenburger A, Behnke S. 2019. Combining simulations and real-robot experiments for Bayesian optimization of bipedal gait stabilization. See Ref. 31, pp. 70–82
89. Zahn B, Fountain J, Houliston T, Biddulph A, Chalup S, Mendes A. 2020. Optimization of robot movements using genetic algorithms and simulation. See Ref. 32, pp. 466–75
90. Iverach-Brereton C, Baltes J, Postnikoff B, Carrier D, Anderson J. 2016. Fuzzy logic control of a humanoid robot on unstable terrain. See Ref. 28, pp. 202–13
91. Böckmann A, Laue T. 2017. Kick motions for the NAO robot using dynamic movement primitives. See Ref. 29, pp. 33–44
92. Seekircher A, Visser U. 2017. A closed-loop gait for humanoid robots combining LIPM with parameter optimization. See Ref. 29, pp. 71–83
93. Masterjohn JG, Polceanu M, Jarrett J, Seekircher A, Buche C, Visser U. 2016. Regression and mental models for decision making on robotic biped goalkeepers. See Ref. 28, pp. 177–89
94. Lanari L, Urbann O, Hutchinson S, Schwarz I. 2017. Boundedness approach to gait planning for the flexible linear inverted pendulum model. See Ref. 29, pp. 58–70
95. Abdolmaleki A, Simões D, Lau N, Reis LP, Neumann G. 2017. Learning a humanoid kick with controlled distance. See Ref. 29, pp. 45–57
96. Kasaei M, Lau N, Pereira A. 2020. A fast and stable omnidirectional walking engine for the Nao humanoid robot. See Ref. 32, pp. 99–111
97. Peña P, Visser U. 2020. Adaptive walk-kick on a bipedal robot. See Ref. 32, pp. 213–26
98. Mitrevski A, Padalkar A, Nguyen M, Ploger PG. 2020. "Lucy, take the noodle box!": domestic object manipulation using movement primitives and whole body motion. See Ref. 32, pp. 189–200
99. Renault B, Saraydaryan J, Simonin O. 2020. Towards S-NAMO: socially-aware navigation among movable obstacles. See Ref. 32, pp. 41–54
100. Makarov PA, Yirtici T, Akkaya N, Aytac E, Say G, et al. 2020. A model-free algorithm of moving ball interception by holonomic robot using geometric approach. See Ref. 32, pp. 166–75
101. Ommer N, Stumpf A, von Stryk O. 2018. Real-time online adaptive feedforward velocity control for unmanned ground vehicles. See Ref. 30, pp. 3–16
102. Balaban D, Fischer A, Biswas J. 2017. A real-time solver for time-optimal control of omnidirectional robots with bounded acceleration. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 8027–32. Piscataway, NJ: IEEE
103. Lobos-Tsunekawa K, Leottau DL, Ruiz-del-Solar J. 2018. Toward real-time decentralized reinforcement learning using finite support basis functions. See Ref. 30, pp. 95–107
104. Watkinson WB, Camp T. 2019. Training a RoboCup striker agent via transferred reinforcement learning. See Ref. 31, pp. 109–21
105. Asada M, Noda S, Tawaratumida S, Hosoda K. 1996. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Mach. Learn.* 23:279–303
106. Abreu M, Reis LP, Lau N. 2020. Learning to run faster in a humanoid robot soccer environment through reinforcement learning. See Ref. 32, pp. 3–15

107. Leiva F, Lobos-Tsunekawa K, Ruiz-del-Solar J. 2020. Collision avoidance for indoor service robots through multimodal deep reinforcement learning. See Ref. 32, pp. 140–53
108. Wiley T, Bratko I, Sammut C. 2018. A machine learning system for controlling a rescue robot. See Ref. 30, pp. 108–19
109. Rizzi C, Johnson CG, Vargas PA. 2018. Fear learning for flexible decision making in RoboCup: a discussion. See Ref. 30, pp. 59–70
110. Simoes D, Lau N, Reis LP. 2019. Adjusted bounded weighted policy learner. See Ref. 31, pp. 324–36
111. Holtz J, Guha A, Biswas J. 2018. Interactive robot transition repair with SMT. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 4905–11. Palo Alto, CA: AAAI Press
112. Celemin C, Ruiz-del-Solar J. 2016. Interactive learning of continuous actions from corrective advice communicated by humans. See Ref. 28, pp. 16–27
113. Grupp M, Kopp P, Huber P, Rätsch M. 2017. A 3D face modelling approach for pose-invariant face recognition in a human-robot environment. See Ref. 29, pp. 121–34
114. Saraydaryan J, Leber R, Jumel F. 2020. People management framework using a 2D camera for human-robot social interactions. See Ref. 32, pp. 268–80
115. Yamaguchi M, Iwamoto G, Abe Y, Tanaka Y, Ishida Y, et al. 2019. Live demonstration: a VLSI implementation of time-domain analog weighted-sum calculation model for intelligent processing on robots. In *2019 IEEE International Symposium on Circuits and Systems*. Piscataway, NJ: IEEE. **https://doi.org/10.1109/ISCAS.2019.8702222**
116. Tanaka Y, Tamukoh H. 2019. Hardware implementation of brain-inspired amygdala model. In *2019 IEEE International Symposium on Circuits and Systems*. Piscataway, NJ: IEEE. **https://doi.org/10.1109/ISCAS.2019.8702430**
117. Lu D, Chen X. 2016. Towards an architecture combining grounding and planning for human-robot interaction. See Ref. 28, pp. 214–25
118. Gemignani G, Veloso M, Nardi D. 2016. Language-based sensing descriptors for robot object grounding. See Ref. 28, pp. 3–15
119. Matamoros M, Harbusch K, Paulus D. 2019. From commands to goal-based dialogs: a roadmap to achieve natural language interaction in RoboCup@Home. See Ref. 31, pp. 217–29
120. Walker N, Peng YT, Cakmak M. 2020. Neural semantic parsing with anonymization for command understanding in general-purpose service robots. See Ref. 32, pp. 337–50
121. Jumel F, Saraydaryan J, Leber R, Matignon L, Lombardi E, et al. 2019. Context aware robot architecture, application to the RoboCup@Home challenge. See Ref. 31, pp. 205–16
122. Peña P, Polceanu M, Lisetti C, Visser U. 2019. eEVA as a real-time multimodal agent human-robot interface. See Ref. 31, pp. 262–74
123. Fukushima T, Nakashima T, Akiyama H. 2019. Mimicking an expert team through the learning of evaluation functions from action sequences. See Ref. 31, pp. 170–80
124. Gabel T, Kloppner P, Godehardt E, Tharwat A. 2019. Communication in soccer simulation: on the use of wiretapping opponent teams. See Ref. 31, pp. 3–15
125. Cooksey P, Mendoza JP, Veloso M. 2017. Opponent-aware ball-manipulation skills for an autonomous soccer robot. See Ref. 29, pp. 84–96
126. Adachi Y, Ito M, Naruse T. 2017. Classifying the strategies of an opponent team based on a sequence of actions in the RoboCup SSL. See Ref. 29, pp. 109–20
127. Schwab D, Zhu Y, Veloso M. 2019. Learning skills for Small Size League RoboCup. See Ref. 31, pp. 83–95
128. Laureano MAP, Tonidandel F. 2020. Analysis of the PSO parameters for a robots positioning system in SSL. See Ref. 32, pp. 126–39
129. Dias R, Cunha B, Azevedo JL, Pereira A, Lau N. 2019. Multi-robot fast-paced coordination with leader election. See Ref. 31, pp. 19–31
130. Dias R, Amaral F, Angelico I, Azevedo JL, Cunha B, et al. 2019. *CAMBADA'2019: team description paper*. Descr. Pap., CAMBADA, Univ. Aveiro, Aveiro, Port. **https://tdp.robocup.org/tdp/2019-tdp-cambada-robocupsoccer-middle-size**

131. Houtman W, Kengan C, van Lith P, ten Berge R, Haverlag M, et al. 2019. *Tech United Eindhoven team description 2019*. Descr. Pap., Tech United Eindhoven, Eindhoven Univ. Technol., Eindhoven, Neth. **https://tdp.robocup.org/tdp/2019-tdp-tech-united-eindhoven-robocupsoccer-middle-size**
132. Wang X, Zhao Y, Chen S, Liu X, Zhang W, et al. 2019. *Water team description 2019*. Descr. Pap., Water, Beijing Inf. Sci. Technol. Univ., Beijing. **https://tdp.robocup.org/tdp/2019-tdp-water-robocupsoccer-middle-size**
133. Yao W, Luo S, Lu H, Xiao J. 2019. Distributed circumnavigation control with dynamic spacing for a heterogeneous multi-robot system. See Ref. 31, pp. 374–86
134. Wong AS, Jeffery R, Turner P, Sleap S, Chalup SK. 2019. RoboCup Junior in the Hunter region: driving the future of robotic STEM education. See Ref. 31, pp. 362–73
135. Hughes J, Shimizu M, Visser A. 2020. A review of robot rescue simulation platforms for robotics education. See Ref. 32, pp. 86–98
136. Zug S, Niemueller T, Hochgeschwender N, Seidensticker K, Seidel M, et al. 2017. An integration challenge to bridge the gap among industry-inspired RoboCup leagues. See Ref. 29, pp. 157–68
137. Gerndt R, Paetzel M, Baltes J, Ly O. 2019. Bridging the gap - on a Humanoid Robotics Rookie League. See Ref. 31, pp. 193–204
138. Pavez M, Ruiz-del-Solar J, Amo V, zu Driehausen FM. 2019. Towards long-term memory for social robots: proposing a new challenge for the RoboCup@Home league. See Ref. 31, pp. 251–61
139. Shimizu M, Takahashi T. 2019. Survey of rescue competitions and proposal of new standard task from ordinary tasks. See Ref. 31, pp. 311–23
140. Gabel T, Falkenberg E, Godehardt E. 2017. Progress in RoboCup revisited: the state of soccer simulation 2D. See Ref. 29, pp. 144–56
141. Michael O, Obst O, Schmidsberger F, Stolzenburg F. 2018. Analysing soccer games with clustering and conceptors. See Ref. 30, pp. 120–31
142. Suzuki Y, Nakashima T. 2020. On the use of simulated future information for evaluating game situations. See Ref. 32, pp. 294–308
143. Michael O, Obst O, Schmidsberger F, Stolzenburg F. 2019. RoboCupSimData: software and data for machine learning from RoboCup Simulation League. See Ref. 31, pp. 230–37
144. Pomas T, Nakashima T. 2019. Evaluation of situations in RoboCup 2D simulations using soccer field images. See Ref. 31, pp. 275–86
145. Suzuki Y, Fukushima T, Thibout L, Nakashima T, Akiyama H. 2020. Game-watching should be more entertaining: real-time application of field-situation prediction to a soccer monitor. See Ref. 32, pp. 439–47
146. Fukushima T, Nakashima T, Akiyama H. 2020. Similarity analysis of action trajectories based on kick distributions. See Ref. 32, pp. 58–70
147. Gabel A, Heuer T, Schiering I, Gerndt R. 2019. Jetson, where is the ball? Using neural networks for ball detection at RoboCup 2017. See Ref. 31, pp. 181–92
148. Speck D, Bestmann M, Barros P. 2019. Towards real-time ball localization using CNNs. See Ref. 31, pp. 337–48
149. Massouh N, Brigato L, Iocchi L. 2020. RoboCup@Home-Objects: benchmarking object recognition for home robots. See Ref. 32, pp. 397–407
150. Kramer ER, Sainz AO, Mitrevski A, Ploger PG. 2020. Tell your robot what to do: evaluation of natural language models for robot command processing. See Ref. 32, pp. 255–67
151. Niemueller T, Reuter S, Ferrein A, Jeschke S, Lakemeyer G. 2016. Evaluation of the RoboCup Logistics League and derived criteria for future competitions. See Ref. 28, pp. 31–43
152. Asada M, Hosoda K, Kuniyoshi Y, Ishiguro H, Inui T, et al. 2009. Cognitive developmental robotics: a survey. *IEEE Trans. Auton. Ment. Dev.* 1:12–34
153. Verschure PF. 2012. Distributed adaptive control: a theory of the mind, brain, body nexus. *Biol. Inspired Cogn. Architect.* 1:55–72