Humanoid Multi-Robot Systems

John E. Anderson, Department of Computer Science University of Manitoba

Synonyms

- Humanoid Multi-Agent Systems
- MAS
- MRS

Related Concepts

- Whole-Body Control of Humanoid Robots
- Whole-body Motion Planning
- Importance of Humanoid Robot Detection
- SLAM and Vision-based Humanoid Navigation
- Multi-Body Simulation
- Human Robot Teaming: Approaches from Joint Action and Dynamical Systems

Definition

Humanoid Robots must not only be able to function autonomously, but also must be able to work effectively as part of a group. Such groups may include other humanoid robots, robots based on different platforms, methods of locomotion, and sizes, and non robot-participants, including software agents and possibly humans. This article reviews major concepts in multi-agent and multi-robot systems and focusses on particular considerations involving teams of humanoid robots.

Background

Teamwork is an important aspect of much human activity. Humans perform activities in groups for two main reasons. The first of these is parallelism and redundancy: a task such as cleaning a floor can be done by a number of humans performing similar operations in tandem. The second is division of skills: in a highly complex environment it may not be possible for all participants to gain knowledge and experience to perform all tasks well. Thus human societies train some individuals in very specialized areas that take many years to learn (e.g. medicine), while others take on particular trades, and jointly a society covers all needs. Similar reasons motivate the use of groups in intelligent systems, and the study of the phenomena surrounding working in groups, including promoting effective teamwork and operating among competitors, falls into the field of *Multi-Agent Systems* (MAS) [18,19]. The main problems associated with MAS include coordinating agents to avoid extraneous work, and promoting coherence, allowing a diverse collection of agents to behave more like a unit [18]. There are many tools used to deal with such problems, including negotiation and authority structures. In a limited sense any real-world problem involving an intelligent agent is a multi-agent system, since the agent itself is generally not the only source of change in the world, but the term is generally used only in situations where there is active involvement with others.

An important distinction in multi-agent systems is whether the system is *homogeneous*, i.e. consisting of identical agents, or *heterogeneous*, supporting distinct types of units [18,19]. Homogeneous systems can be simpler in that they allow agents to make the assumption that all others they work with are like them, as well as being simpler to construct, since every agent is physically the same and the range of skills is identical for each agent. Bounded agent resources and complex domains, however, point toward heterogeneity: having specialized skills for different agents, and specialized equipment or tools to support those skills. This is seen in many human domains, where people specialize in particular jobs. Practical considerations of expense similarly point to heterogeneity. Some skills may require very expensive tools or equipment, or knowledge that is expensive to acquire, and equipping each agent with these may be cost-prohibitive. Parsimony in design and implementation also means not including skills in every agent when these may only be used infrequently and provided by other specialized agents.

Heterogeneity is also practically necessary in domains where agents can be damaged or destroyed [13,17]. In such scenarios, the likelihood of destruction further emphasizes economic arguments for heterogeneity: cheaper, simpler agents can be exposed to higher levels of risk, preserving those with expensive or rarer skills.

Thus far, the term *agent* has been used exclusively, and work in MAS utilizes a wide range of abstract agents, from individual participants in economic theories to software agents in social media. Work that is grounded in the physical world in the form of groups of robots is referred to as *Multi-Robot Systems* (MRS). Multi-robot systems differ from other work in MAS in that they are necessarily grounded in physical space and must obey the restrictions of the physical world. Abstract agents can move without considering real world error in perceived location, perceive the world without sensor error, and occupy the same space, while control techniques for robots are significantly impacted by all of these elements and many more. Much work in Multi-Agent Systems purports to be transferable to groups of robots, but without a physical implementation demonstrating this, such transferability is questionable simply because issues involving physical space, sensor error, etc. are limited in consideration compared to testing in the physical world. Compare, for example, the range of papers in modern conferences such as AAMAS [1] with the range of systems that are actually demonstrated at conferences in robot embodiment.

Moving to Multi-Robot Systems brings the physical world to bear on issues such as coordination and cohesion. In a cleaning scenario with a group of robots, for example, we are limited in the number of robotic cleaners that can be added by the size of the space available, and well before that space is completely filled, the amount of interference between robots will outweigh the benefits of adding more [12]). Knowledge communicated between robots will have issues of timeliness and error because of the nature of physical communication. In domains where robots can be lost, damaged, or destroyed, this also means that knowledge of the current members of a team will differ between individuals, affecting coordination mechanisms [13,17].

Humanoid MRS: Theory and Application

Until recent years, with a few significant exceptions, humanoid robots have not figured prominently in multi-robot systems work. Important early MRS work involving physical robots tended to restrict platforms to fairly simple (by today's standards) wheeled and treaded models with limited perception (e.g. [3]). Even later work tended to avoid complex forms such as humanoids in favour of more basic wheeled platforms. (e.g. [10]). At the time there were very good reasons for this: good control models for wheeled platforms had already existed for some time, while researchers were still experimenting with basic walking gaits for humanoids. Trying to control enough motors to provide 20 Degrees of Freedom (DOF) on a computational platform that was portable and low-powered enough for a humanoid was a significant challenge, and so using humanoids was a major impediment to those interested primarily in, for example, multi-robot simultaneous localization and mapping (SLAM), coordinated exploration, and foraging. Simply keeping a population robots running at all for repeatable long-term experiments was very difficult, leading to the development of better robot simulators for such experiments. Other practical concerns included the portability of batteries and computing devices to allow a humanoid robot to be mobile and run autonomously for a reasonable length of time. Beyond these issues, simply affording more than one or two humanoid robots was a real challenge for most laboratories since commercial humanoid robots were generally not available.

The one area where humanoids were employed in multi-robot settings from a comparatively early date was in robotics competitions involving soccer, specifically geared to humanoids because the game itself was played by humans and intended to form a grand research challenge. Both RoboCup and FIRA debuted humanoid robot competitions in 2002 [7,8]. These began with demonstrations of soccer-based skills such as penalty kicks [11] and in RoboCup, gravitated toward robotic soccer games as a standard challenge problem as soon as technology was sufficiently mature to support this (2-on-2 soccer was first played at RoboCup in 2005 [7], the first time the task required potential teamwork). In FIRA, through the HuroCup, the focus remained on a breadth of events to test humanoid robotic skills, analogous to human multi-event athletic competitions such as decathlons. While soccer has at times been among these it is one event amid many others challenging the same robots, including obstacle courses, basketball, climbing, and archery [8].

Robotic soccer was used as a challenge problem in RoboCup in part because of the multi-robot aspects of soccer: a collection of robots that do not coordinate as a team can do only so well in soccer against those that do coordinate. Thus over the years, a range of results has been produced supporting teamwork in soccer applications, that are intended to apply to humanoid multi-robot systems in general. These range from adapting individual control to better operate in the presence of others, such as recovering after bumping into other players [16], or recognizing opponents and teammates on the field [14], to schemes for role assignment on the field, coordinating plays, etc. [9].

This work has also helped to further an understanding of robotic soccer beyond humanoids (since elements of coordination, for example, can be adapted to other types of robots), and has also led to direct work outside of soccer (for example, work on detecting other humanoid robots is directly applicable to detecting pedestrians, an important part of work toward autonomous cars). Work on robotic soccer has also led to advances in mechatronics that in turn have led to better humanoid robots. Humanoid robotic platforms have been enhanced in recent years by the production of commercial platforms that have begun to be produced in numbers large enough to become reasonably affordable [11]. Some of these, such as the DARwIn-OP, have basic soccer skills built in, allowing individuals or labs to begin orienting work toward competitive soccer teams more quickly.

As humanoid robotic technology improved, and the skills necessary to function well in more controlled competition domains such as soccer (e.g. SLAM is physically easier in Robotic soccer because of the presence of visually obvious lines on the field [4]), more researchers began to apply what they learned from competitions to other domains. For example, SLAM provides interesting challenges for humanoid robots compared to wheeled platforms, because the motion model of humanoids is much more complicated [2,6]. A twisted body pose, for example, can provide a similar visual view as facing a different direction, and while wheeled robots accumulate error in positioning based on wheel slippage when moving, motion across many joints and the effects of step differences on uneven terrain can greatly magnify this in humanoids. Moreover, vision-based SLAM in high-DOF humanoid robots must account for head motion and the fact that the camera view is not stable as the robot takes steps, but shifts at various points in a walking gait. Baltes et al., using robots and basic skills developed for both the RoboCup humanoid competition and the FIRA HuroCup, developed an approach to multi-robot SLAM dealing with the volume of error accumulated under these conditions with small humanoid robots [2,6]. This was designed to operate under low computational resources (running on mobile phones carried by the robots that were also running all robotic and real time vision code (Figure 1). Inspired by the HuroCup obstacle run, multiple robots mapped an area walled off with colored obstacles, and gradually constructed a shared map by communicating partial findings over time and merging a shared map.

From the perspective of humanoid multi-robot systems, there are still strong elements of homogeneity in most modern robotic soccer teams. Competitions typically stick to smaller numbers (e.g. 3-on-3, 4-on-4) rather than the standard 11 in human soccer. This is partly due to the expense of maintaining (and shipping) large numbers of robots, but is also driven by a competing desire for better humanoid robots, which in turn increase the expense involved. As

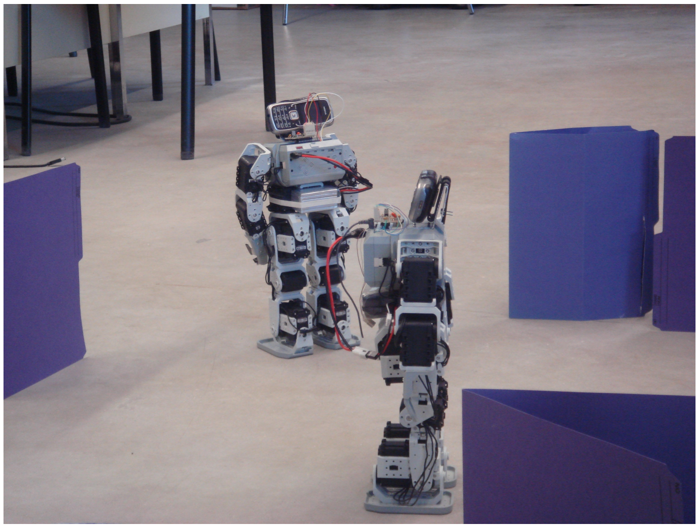

Fig. 1. Multi-Robot Simultaneous Localization and Mapping on Small Humanoid Robots

robots and peripheral hardware such as power supplies become more advanced, it becomes feasible and desirable to move to larger human forms that better approximate the humans, and so there is more focus on "teen" (mid-size) and adult size leagues [11] where robots are considerably more expensive to develop and maintain. In terms of heterogeneity, this leads to a smaller number of distinct playing roles than might be seen in many human activities: at worst, a unique playing style for a goalkeeper vs other players that function as strikers and cover dynamically changing areas of the field. Even one distinct role, or one distinct set of code, will technically make a system heterogeneous, but this is a less diverse set of skills than would be seen in social structures more advanced than a soccer team.

In recent years the HuroCup has made changes to the standard soccer event to both encourage further work in humanoid multi-robot systems and increase heterogeneity in soccer. The HuroCup organizers developed a *United Soccer* event [11], where robots from the range of entrants were placed on one of two soccer teams randomly, forcing robots developed separately to work with one another without prior knowledge. A common communications protocol is provided, along with a basic ontology for communicating intentions and the state of the game. Beyond uncertainty of the abilities of other players, this magnifies the problem in terms of heterogeneity since all robots run distinct code, and within the size and proportion limits they can all be physically distinct from one another as well. This changes not only the speed at which robots move, but how they physically appear, making passing much more challenging since one must recognize a broader range of robots visually. A separate motivation for united soccer was encouraging work with larger team sizes despite many labs not being able to afford a significant number of humanoid robots [11]. Because of differences in budgets and personnel involved, the team entries (like all competitions) are in a range of qualities beyond the minimum for qualification. The presence of a mix of strong soccer players along with weaker members makes functioning as a team much more challenging as well. Because of this, individual points are tallied for robots for various actions such as successful passes, in addition to regular soccer play. In order to provide point opportunities for weaker robots as well, and further this event as a research challenge, the field is modified to include boxes that can be used as passing targets for alternative points in case weaker robots cannot recognize teammates (Figure 2).

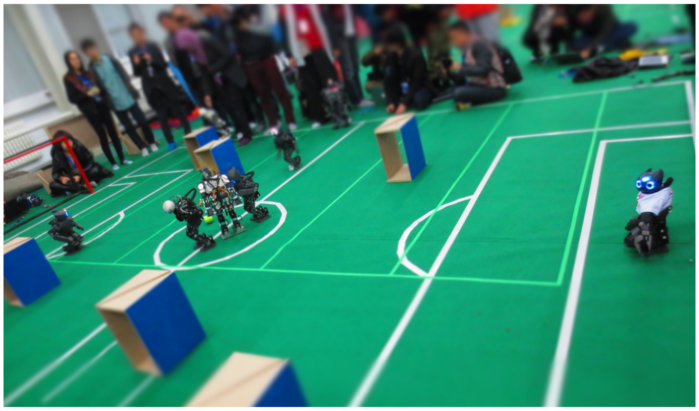

Fig. 2. HuroCup United Soccer Event, Beijing, 2014

The ability to have reasonably-priced humanoid robots has led some in the RoboCup and FIRA communities to also do work beyond soccer using the equipment employed in robotic competitions each year, and some of this work focusses on more heterogeneous scenarios. Notably, Keiner and von Stryk [15] performed interesting work on heterogeneous task allocation in multi-robot systems, using one of the humanoid robots they originally employed at the RoboCup competitions. Their work was one of the first instances attempting to physically implement a heterogeneous multi-robot system consisting of distinct robotic platforms including humanoids, with each platform being used according to its strengths in particular tasks. The scenario they implemented was the need to put a ball into a target goal that was further than the distance a humanoid could walk on a typical battery charge. In addition to the humanoid, a wheeled platform able to carry the humanoid was included in the system, and the abilities of each robot type (e.g. humanoids can kick a ball, a wheeled platform can travel long distances, humanoids can climb onto wheeled platforms) were expressed as part of a task allocation system that could properly allocate tasks to the unit best suited to carry them out. In their test scenario, tasks were broken down, assigned to the most appropriate robot, and most importantly properly timed and coupled so that the wheeled robot would position itself to be convenient to the humanoid, who then boarded the wheeled robot, was carried across a distance to a place convenient to a ball, and then dismounted and properly kicked the ball. From an abstract multi-agent systems perspective the task allocation was not particularly challenging (two agents of very distinct types and little confusion as to which should likely receive each task). However, from a multi-robot systems perspective, the fine integration between tasks allocated to different types of robots, and their understanding of one another (e.g. properly positioning the wheeled robot and timing this to the humanoid seating itself safely) was very sophisticated and beyond most abstract work in simulation.

Open problems

There are many open problems in humanoid multi-robot systems. Many of the conventional problems associated with working with multiple robots (forming coalitions, dividing tasks, etc.) are multiplied in difficulty by the greater variability of the humanoid form, and some well-known problems must be reexamined for sophisticated humanoids. For example, work in SLAM, including multi-robot SLAM, has progressed far enough that deployed solutions for automated factories are possible, and this area is thought to be fairly mature. However, applying these approaches to humanoid robots means greatly increasing the sophistication of motion control models, and as humanoid robots become more and more sophisticated (e.g. even a 20 DOF humanoid, common in today's robotic soccer competitions, is highly constrained compared to a real humanoid), so will these models, thus directly affecting the functionality of SLAM approaches. The enormous range of possible motions that must be considered when doing full-body motion planning in humanoids (e.g. [5]) illustrates the similar range in possible errors in pose that have to be considered when performing SLAM. Thus we still see SLAM approaches, even those claiming to be useful on humanoids, demonstrated using robots with wheels rather than a true walking gait (e.g. [**?**], which uses a human torso mounted on wheels).

Other problems arise from the significant expansion in the concept of heterogeneity that the humanoid form brings to multi-robot systems. While some of the issues involved have been pointed out earlier, in terms of recognizing other humanoids in groups when they vary in size and shape, and dealing with different skill sets and abilities in soccer, this is really only the tip of the iceberg in terms of humanoid variability. In the real world, humans continue to function despite fairly significant damage (e.g. walking with a limp because of a damaged muscle or joint). In current humanoid robotic systems, robots undergoing degrees of malfunction or damage are generally removed from the application for repair or replacement (e.g. in robotic soccer play, humanoids often lose functionality on the field because of dead batteries, broken servos, etc., and are then removed from the field). In many potential robotic domains where groups of humanoid robots could be placed, such losses would be unacceptable because of expense or difficulty replacing robots (e.g. space exploration, search and rescue, battlefield scenarios) and robots would be expected to be leveraged to the best of their abilities. In such scenarios, we have the potential not only for different sizes and appearances, but huge changes in range of motion and functionality over time that must be considered as part of a multi-robot system. A humanoid with a broken leg could still crawl, and be of use to its team, and these differences could be adapted to through the support of other robots (e.g. leading a blind teammate that can still perform other operations, or helping to support another that has balancing difficulties because of a bad knee). While there are schemes for adapting coalitions and changing task allocation based on changing skill sets, the range of distinct robot types over time is combinatorially larger with high-DOF humanoids compared to simpler wheeled platforms (e.g. a partial failure in one of 20+ DOF as opposed to one of two wheels damaged). Current successes in humanoid multi-robot systems have relied heavily on better and more affordable humanoid robots, and further progression along these lines will similarly require more maturity and further successes in humanoid robots themselves.

Recommended Readings

[1] AAMAS 2017 conference. `http://www.aamas2017.org/`. Accessed: 2017-07-31.

[2] Jonathan Bagot, John Anderson, and Jacky Baltes. Vision-based multi-agent slam for humanoid robots. In *Proceedings of the 5th International Conference on Computational Intelligence, Robotics and Autonomous Systems (CIRAS-2008)*, pages 171–176, June 2008.

[3] Tucker Balch. The impact of diversity on performance in multi-robot foraging. In *Proceedings of the Third Annual Conference on Autonomous Agents*, AGENTS '99, pages 92–99, New York, NY, USA, 1999. ACM.

[4] Jacky Baltes. Localization for mobile robots using lines. In *Proceedings of the Seventh International Conference on Control, Automation, Robotics and Vision (ICARCV)*, December 2002.

[5] Jacky Baltes, Jonathan Bagot, Soroush Sadeghnejad, John Anderson, and Chen-Hsien Hsu. Full-body motion planning for humanoid robots using rapidly exploring random trees. *KI - Künstliche Intelligenz*, pages 1–11, 2016.

[6] Jacky Baltes, Chi Tai Cheng, Jonathan Bagot, and John Anderson. Vision-based obstacle run for teams of humanoid robots (demonstrated system). In *Proceedings of the 10th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2011)*, pages 1319–1320, Taipei, Taiwan, May 2011.

[7] Jacky Baltes, N. Michael Mayer, John Anderson, Kuo-Yang Tu, and Alan Liu. The humanoid leagues in robot soccer competitions. In *Proceedings of the IJCAI Workshop on Competitions in Artificial Intelligence and Robotics*, pages 9–16, Pasadena, California, July 2009. AAAI Press.

[8] Jacky Baltes, Kuo-Yang Tu, Soroush Sadeghnejad, and John Anderson. Hurocup: Competition for multi-event humanoid robot athletes. *The Knowledge Engineering Review*, FirstView:1–14, August 2016.

[9] Sven Behnke and Jörg Stückler. Hierarchical reacive control for humanoid soccer robots. *International Journal of Humanoid Robotics*, 5:375–396, 2008.

[10] Wolfram Burgard, Mark Moorsy, Cyrill Stachniss, and Frank Schneidery. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376–386, 2005.

[11] Reinhard Gerndt, Daniel Seifert, Jacky Baltes, Soroush Sadeghnejad, and Sven Behnke. Humanoid robots in soccer - robots versus humans in robocup 2050. *IEEE-RAS Robotics and Automation Magazine*, 22(3):147–154, September 2015.

[12] Dani Goldberg and Maja J. Matarić. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence*, AAAI'97/IAAI'97, pages 637–642. AAAI Press, 1997.

[13] Tyler Gunn and John Anderson. Dynamic heterogeneous team formation for robotic urban search and rescue. *Journal of Computer and System Sciences*, 81(3):553–567, May 2015.

[14] Mohammad Javadi, Sina Mokhtarzadeh Azar, Sajjad Azami, Saeed Shiry Ghidary, Soroush Sadeghnejad, and Jacky Baltes. Humanoid robot detection using deep learning: A speed-accuracy tradeoff. In *Proceedings of RoboCup-2017*, Nagoya, Japan, July 2017.

[15] Jutta Kiener and Oskar Von Stryk. Cooperation of heterogeneous, autonomous robots: A case study of humanoid and wheeled robots. In *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2007)*, pages 959–964. IEEE, 2007.

[16] Marcell Missura, Cedrick Münstermann, Philipp Allgeuer, Max Schwarz, Julio Pastrana, Sebastian Schueller, Michael Schreiber, and Sven Behnke. Learning to improve capture steps for disturbance rejection in humanoid soccer. In *Proceedings of RoboCup-2013*, Eindhoven, Netherlands, July 2013.

[17] Geoff Nagy and John Anderson. Active team management strategies for multi-robot teams in dangerous environments. In *Advances in Artificial Intelligence: 30th Canadian Conference on Artificial Intelligence*, pages 385–396, Edmonton, AB, May 2017. Best Paper Award.

[18] Gerhard Weiss, editor. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence.* MIT Press, Cambridge, MA, USA, 1999.

[19] Michael Wooldridge. *An Introduction to Multiagent Systems.* John Wiley & Sons, Inc., New York, NY, USA, second edition, 2009.