



This is a repository copy of *Multi-robot systems research: a data-driven trend analysis*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/201044/>

Version: Accepted Version

---

**Proceedings Paper:**

Amorim Marques, J.V., Lorente, M.-T. and Gross, R. [orcid.org/0000-0003-1826-1375](https://orcid.org/0000-0003-1826-1375)  
(2024) Multi-robot systems research: a data-driven trend analysis. In: Siciliano, B. and Khatib, O., (eds.) Distributed Autonomous Robotic Systems. DARS 22. 16th International Symposium on Distributed Autonomous Robotic Systems 2022 (DARS 2022), 28-30 Nov 2022, Montbéliard, France. Springer Tracts in Advanced Robotics, 28 . Springer , pp. 537-549. ISBN 978-3-031-51496-8

[https://doi.org/10.1007/978-3-031-51497-5\\_38](https://doi.org/10.1007/978-3-031-51497-5_38)

---

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: [http://dx.doi.org/10.1007/978-3-031-51497-5\\_38](http://dx.doi.org/10.1007/978-3-031-51497-5_38)

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Multi-Robot Systems Research: A Data-Driven Trend Analysis

João V. Amorim Marques<sup>1,2</sup>, María-Teresa Lorente<sup>1,3</sup>, and Roderich Groß<sup>1\*</sup>

<sup>1</sup> Sheffield Robotics, The University of Sheffield, Sheffield, UK,  
`r.gross@sheffield.ac.uk`

<sup>2</sup> INESC Science & Technology, Faculty of Engineering, University of Porto, Portugal

<sup>3</sup> Instituto Tecnológico de Aragón, Zaragoza, Spain

**Abstract.** This paper provides a data-driven analysis of the trends of research on multi-robot systems (MRS). First, it reports the findings of an exhaustive search of the MRS studies published from 2010 to 2020 in 27 leading robotics journals, including a quantitative analysis of trends. Second, it reports the findings of a survey capturing the views of 68 leading experts in the field of MRSs. Finally, it summarises the findings.

## 1 Introduction

In the late 1940's, Walter [1] built two autonomous robot tortoises that displayed not only impressive individual capabilities, including moving towards a light source or into a station for self-charging, but also collective capabilities, described by the author as “mutual recognition” or forming a “sort of community”. These robots are one of the first physical realisations of multi-robot systems (MRSs). Despite significant developments, especially over the past 35 years, MRSs are only beginning to become commonplace in real-world applications. Numerous surveys have been published on the topic of MRSs [2,3,4,5], and the reader is referred to these for a comprehensive overview of the field. This paper seeks to complement previous studies by providing the first data-driven analysis of the trends in the field. It reports:

1. the quantitative findings of an *exhaustive search* of the MRS studies published in leading robotics journals (in between 2010 and 2020); and
2. an analysis of the views of leading MRS experts.

The findings may provide insights to guide future research, for example, by determining areas of development that require significant effort and the short-term to long-term prospects for specific applications.

## 2 Exhaustive Search

This section presents the methods and findings from the exhaustive search.

---

\*All authors contributed equally to this work.

## 2.1 Methodology

We use *robot* to refer to a machine that perceives, and acts in, an environment in physically plausible ways. The *environment* can be understood as the set of characteristics that define the space in which the robot resides. Studies employing abstract environment/robot models are considered within scope where clearly of relevance to real-world MRSs. For the purpose of the exhaustive search, we restrict the scope to systems composed of mobile robots, that is, robots that are able to change their position and orientation in their environment.


A *multi-robot system* can be defined as a group of robots that operate in the same environment while exhibiting some form of interaction. When part of an MRS, robots usually work collaboratively towards a common goal. This allows them to address a more complex task, or perform a task more efficiently, than when on their own. MRSs may have different levels of autonomy when deployed, from being teleoperated to being fully autonomous.

We use an exhaustive search to target all documents published in years 2010–2020 in 27 of the 28 “robotics” journals that are included in the 2020 Journal Citations Report of the Web of Science Group.<sup>4</sup> The documents include original research articles, perspectives, and surveys, but exclude editorials, corrections and other communications. Conference papers, although valuable to the robotics field, were not considered due to the sheer volume of works and the lack of clear criteria for which conferences to consider. For each document, we determine whether it is within scope. Where a document is not within scope, no further analysis is performed. For each document within scope, it is first rated in terms of (i) the originality (up to 6 points) of the work that is reported, (ii) its rigour (up to 4 points) and (iii) its significance (up to 5 points). The individual ratings are then added to produce a *relevance* score. Although detailed criteria and training are used to guide the assessors (two of the authors) and improve consistency in evaluation, the ratings are finally subjective. To reduce the impact of possible assessor bias (i.e. one assessor tending to provide higher values than the other), samples of documents were marked by both of the assessors—once at the beginning and twice during the exhaustive search. From the results of these double-marking rounds, a consistent but small bias can be observed. To counter this bias, the relevance scores that we report have been adjusted by -0.15 and +0.20 respectively, such that the top 25% of documents within scope for either assessor have a score of 6.00 or above. In the following, these works are referred to as highly relevant.

## 2.2 Findings from journal-specific analysis


Overall, 24 250 documents were examined in the exhaustive search. Of these, 1487 (6.1%) are within scope, and hence analysed in detail. Figure 1 shows the percentage of documents that are within scope for each journal. For eight

<sup>4</sup>The documents of one of the journals (International Journal of Robotics & Automation) could not be accessed in full and were hence excluded.




**Fig. 1.** Percentage of documents within scope (per journal). Portions in red correspond to papers regarded as highly relevant. The numbers represent the total number of papers identified to be within scope.

journals, more than 10% of documents are within scope. The journals Swarm Intelligence (SI) and Autonomous Robots (AURO) have the strongest focus on MRS research.



**Fig. 2.** Distribution of the technology readiness levels for the documents within scope (per journal).

Figure 2 shows the distribution of the technology readiness levels (TRL) that we estimated for the systems reported in all documents within scope for each journal (for a definition, see [6]). Most of the journals primarily report findings with a TRL level of up to 4, considered as laboratory validation by the assessors. The Journal of Field Robotics (JFR) and Revista Iberoamericana de Automática e Informática Industrial (RIAI) stand out: 75% and 46% of documents report work of TRL level 5 or above, owing to these journals primarily focusing on studies that have been validation via field experiments. It should be noted that journals RIAI, Intelligent Service Robotics (ISR), JFR, The International Journal of Robotics Research (IJRR), AURO and Robotica introduce



**Fig. 3.** Distribution of documents within scope by type of validation (laboratory only, real-world only, or both).

some documents evaluated with a TRL of 6, where the system was demonstrated in a relevant environment.


Figure 3 shows the percentage of documents within scope that validated their research via laboratory and/or real-world experiments, as well as the percentage of documents that report formal, theoretical results. As expected, the JFR presents the greatest proportion of documents with real-world validation. In terms of theoretical validation, the journals IEEE Transactions on Robotics (T-RO), IJRR, AURO, IEEE Robotics and Automation Letters (RA-L), Robotica, and Journal of Intelligent & Robotic Systems (JINT) contain the largest proportions of works reporting theoretical findings (journals with fewer than 10 works being omitted from consideration of proportions).

### 2.3 Findings from trend analysis


This section presents a quantitative analysis of the trends observed based on the data collected in the exhaustive search.

Figure 4 presents the number and percentage of documents that were within scope for each year. Dashed lines show the trends, obtained through linear regression using the least squares method. The trend in number of documents on MRSs suggests a substantial growth in the number of documents within scope, whereas the trend in percentage of documents suggests that the fraction of documents devoted to MRSs has decreased. A possible explanation is that robotics journals tend to report on an increasingly diverse range of topics, some of which attracted significant attention in recent years (e.g. soft-bodied and learning robots).

Figure 5 presents the evolution of the percentage of studies within scope that perform theoretical, laboratory and real-world validation. Note that it is possible for a single study to feature for multiple types of validation. Overall, the trends suggest the percentage of theoretical studies is increasing, while laboratory validation is shown to be decreasing and real-world validation is shown to remain relatively constant over the years. Note however that the number of documents



**Fig. 4.** Number and percentage of documents within scope (per year), as identified by the exhaustive search. The dashed lines represent the trends obtained by linear regression.



**Fig. 5.** Percentage of documents within scope using theoretical, laboratory and real-world validation (per year), as identified by the exhaustive search. The dashed lines represent the trends obtained by linear regression.

within scope is increasing over the years. As a result, the total amount of validation is increasing as well, for both laboratory and real-world experiments. Although theoretical validation is often based on strong assumptions, it can offer a detailed understanding of the underlying mechanisms and allow for a basic demonstration of the validity of the proposed solutions. The decreased focus on laboratory validation may be attributed to increases in computational capabilities, allowing more trustworthy and realistic simulations to be performed, which may help substitute for laboratory validations.

### 3 Expert Survey

The expert survey is conducted to help better understand the state of the art and possible future developments of MRSs. It captures the views of a representative sample of internationally leading experts in the field of MRSs.

### 3.1 Methodology

The expert survey received ethics approval following a review by The University of Sheffield (application reference 034028). It uses a questionnaire that includes closed- and open-ended questions. Responses are anonymous. Participant can choose not to answer any specific question if they prefer. Where subject-specific questions provide multiple options to choose from, these are presented in a random order to each participant.

To identify suitable prospective participants, research paper databases; conference, workshop and journal websites (invited speakers, organisers, editorial boards); resources related to non-academic organisations (e.g. industry, government); websites of technical committees (e.g. IEEE RAS TC Multi-Robot Systems); and other sources are explored. Special emphasis is placed on identifying female experts, experts from traditionally underrepresented geographic regions, and experts from industry and other non-academic organisations. The participants are given three weeks to complete the survey from the moment of invitation (9 July 2020).

### 3.2 Findings

Overall, 179 experts were invited to participate in the survey, including 119 men (66.5%) and 60 women (33.5%).<sup>5</sup> The invited experts were from the following broad, geographic regions: Africa (10), America (62), Asia (35), Europe (66), and Oceania (6), as defined by the Statistics Division of the United Nations<sup>6</sup>.

A total of 68 participants completed the survey, corresponding to a participation of 38.0% of the invited experts. Each participant provided consent to participate in the survey. In this section, we analyse the findings.

**Participant distribution** To ensure the responses to the expert survey are representative, or reveal biases where they exist, we included four personal questions regarding the participant gender, geographic region, sector of work, and the number of years of relevant professional experience. Overall, 76.5% and 20.6% of participants reported themselves as male and female, respectively. A small percentage of participants preferred not to answer (2.9%). The gender distribution is clearly biased towards men, primarily as 66.5% of those invited to participate were men. This can, at least partially, be attributed to women being underrepresented in the field of MRSs, and, especially, in leading roles. For the related disciplines of electrical engineering, mechanical engineering, and computer science, in between 13.5% and 19.7% of tenured/tenure-track faculty positions in the USA are held by women according to a recent survey [7].

Most participants were from Europe, America and Asia. The participant share of Europe is 42.6%, with 1.5% from Eastern Europe, 7.4% from Northern

<sup>5</sup>The gender for each prospective participant was estimated. Actual participants could choose to report their gender.

<sup>6</sup>See geographic regions on <https://unstats.un.org/unsd/methodology/m49/>

Europe, 16.2% from Southern Europe, and 17.6% from Western Europe. The participant share of America is 32.4%, with 26.5% from North America and 5.9% from South America. The participant share of Asia is 20.6%, with 7.4% from Eastern Asia and 4.4% from each of South-eastern, Southern, or Western Asia. The remaining participants were from Australia and New-Zealand (2.9%) or did not disclose their geographic region (1.5%). Although the survey succeeded in capturing the views of experts from a wide range of regions, some bias can be observed, in particular towards participants from Europe. It should also be noted that none of the experts reported Africa as their geographic region.


Most participants work in academia (77.9%), followed by industry (14.7%). Participants who work in both academia and industry, or in other sectors, are in the minority (2.9% for either case).

Three quarters of the participants have up to 20 years of relevant professional experience with half reporting 10 to 20 years (50.0%) and a quarter less than 10 years (25.0%). Most of the remaining participants have between 20 and 30 years of relevant professional experience (19.1%). Two participants have more than 30 years of relevant professional experience (2.9%).

**Analysis of responses to close-ended questions** This section presents the views obtained from the experts when answering the subject-specific, closed-ended questions of the survey.

The first question is about what constitutes an MRS. The experts are asked to what extent they consider 12 example systems, labelled  $a$ – $\ell$ , to be MRSs. Figure 6 shows the distribution of their responses. Their opinions are diverse. However, broad consensus exists that systems  $a$ ,  $b$  and  $c$  would not be considered as MRSs. In these systems, the robots may not even be active. Moreover, broad consensus exists that system  $\ell$ , where the robots operate in a same environment and have a common purpose, would be considered as an MRS. Most experts viewed systems where multiple robots interact, such as systems  $i$ ,  $j$ , and  $k$ , as instances of MRSs, even when the robots were physically separated, such as in system  $k$ . Systems  $d$  and  $e$  are not considered as MRSs by most of the experts, though conflicting views exist. For system  $d$ , the robots are not self-propelled and do not necessarily interact. For system  $e$ , one robot transports another robot. For the remaining systems  $f$ – $h$ , no clear trend is observed. For system  $f$ , the robots share a common infrastructure to obtain position information but are otherwise independent. In the case of system  $g$ , a group of conventional manipulators work towards a common goal along a factory assembly line. Given that the motion of each manipulator is fully automated, there appears to be limited interaction. Where the operations of one manipulator builds upon the operations of the previous ones, however, the group could be seen as interacting via the products that they assemble. For system  $h$ , each robot is teleoperated by a dedicated human operator while sharing the same environment. Different from system  $j$ , the robots are not moving autonomously. As the wider context may not be sufficiently clear, it is difficult to categorise this system, as suggested






**Fig. 6.** Distribution of the expert responses to the question of whether the following are considered an MRS: (a) A group of robots stored on the same shelf, (b) Any robot of which more than two copies exist, (c) A group of robots that can be stacked to ease transportation, (d) A group of passive beads residing in a same environment, and being externally controlled through changes in a magnetic field, (e) A robot transporting another robot, (f) A group of robots operating in different environments, and that do not interact; each obtains positional information from the same global infrastructure, (g) A group of conventional manipulators along a factory assembly line, working on the same products though at different stages, (h) A group of robots operating in a same environment, each being teleoperated by a dedicated operator; the operators may have conflicting goals, though they always attempt to prevent collisions between robots, (i) A robot composed of several modules that can autonomously rearrange the way they are connected to assume a new shape, (j) A group of robots operating in a same environment, each having its own goal, including not to collide with others, (k) A group of robots operating in different environments and that learn from each other through a cloud service, (l) A group of robots operating in a same environment and having a common purpose. Statements ranked by weighted mean results.

by a relatively high percentage of experts selecting the option “other (e.g. would need more information to decide)”.

The second question concerns the future deployments of MRSs. The experts were asked the time frame in which they expect MRSs to be widely deployed within their geographic region in various application scenarios. Figure 7 illustrates the distribution of the expected times for each of the application scenarios. Almost all experts consider warehouse logistics as an achievable target in the short term (0–5 years). The majority of experts consider surveillance and security, manufacturing and entertainment as achievable targets in the short term (0–5 years). These application scenarios have already fully embraced robotics in multiple geographic regions. Transport and delivery, inspection and maintenance, agriculture and forestry, education, mining, first response, and domestic applications are considered achievable targets in the medium term (5–10 years) by many of the experts. Space, medical applications, rehabilitation and care and construction are expected to face a longer wait (> 10 years). Additionally, we should notice that education has a moderate number of experts suggesting MRSs may never be widely deployed or not be applicable.



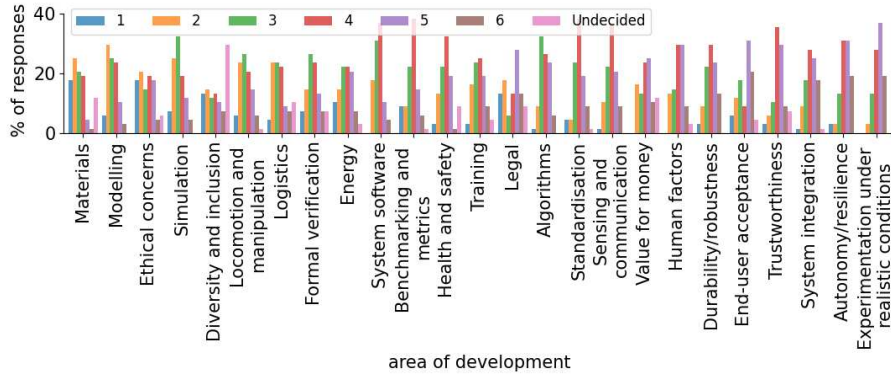
**Fig. 7.** Distribution of the responses by experts to the question in which time frame they expect MRSs to be widely deployed within their geographic region in the specified applications. Applications ranked by weighted mean results.

Deploying MRSs in the real world involves solving some problems that are equally present when deploying single robot systems. The third question specifically focuses on the transition from systems of single robots to systems of multiple robots. It asks the expert to gauge how much *additional* effort this transition requires by the robotics community with respect to a number of areas of development. Figure 8 shows the distribution of the additional effort the experts consider necessary for each of the suggested areas of development. In general, we observe that all aspects would demand additional effort. Particularly, experimentation under realistic conditions, autonomy/resilience, system integration, trustworthiness, and end-user acceptance can be highlighted as the most difficult areas of development, with the highest percentage of experts choosing ratings of either 5 or 6. On the contrary, issues such as materials, modelling, ethical concerns, and simulation are considered less demanding in terms of additional effort, with the highest percentage of experts choosing ratings of either 1 or 2. Regarding legal aspects, the opinions are divided. Diversity and inclusion has been appraised as “Undecided” by 29.4% of experts, possibly as the challenge is viewed by some as the same irrespective of whether the system to be designed comprises one or multiple robots.


The last question asks the experts how fast they expect the field of MRSs to have grown in ten years from the time of the survey compared to the field of robotics as a whole. Figure 9 presents the responses, suggesting a bimodal distribution with a small peak at “slower” and a larger one at “faster”.

**Analysis of responses to open-ended questions** In this section, we present and discuss the participant responses to the open-ended questions of the survey.

The first question asked how the experts attitude towards MRSs had changed over the past 10 years. About a third (32.4%) of the participants responded that



**Fig. 8.** Distribution of the expert responses to the question of how much additional effort is required by the robotics community to move from single-robot system to MRS in real-world applications with respect to various areas of development. The scale is such that 1 corresponds to “insignificant”, whereas 6 corresponds to “unprecedented (e.g. paradigm change required)”. Areas ranked by weighted mean results.



**Fig. 9.** Distribution of the expert responses to the question of how fast they expect the field of MRSs to have grown in ten years from the time of the survey compared to the field of robotics as a whole.

their attitude had not changed. The remaining participants presented a multitude of aspects that had changed when considering their position towards MRSs. The most frequently mentioned aspects include the focus of their work (27.3%), in particular a shift from theoretical work to more practical applications, the definition of MRS they used, to either include or exclude specific control architectures (centralised vs. decentralised), system size, etc., and the consideration of new capabilities or applications to MRSs, such as machine learning, cloud computing, and warehouse logistics. Although the outlook of most responses to this question were positive, a small number of participants (6.8% of the participants that reported a change in attitude) reported a negative position, questioning the extent to which real-world applications will be impacted on by MRSs.

The next question asked participants to give examples of commercial applications other than warehouse logistics where MRSs are currently being deployed. A total of 63.2% participants gave examples, with the most common being (i)

entertainment, for example, through the use of drones such as in light shows and for filming, (ii) agriculture, for seeding, tending and harvesting of crops, and (iii) surveillance. Interestingly, agriculture is included in the list of applications that experts expect to reach deployment only in the medium term (see Section 3.2). Other applications, from the most to the least mentioned, include defence, inspection, monitoring, logistics, transport, education, hospital applications, mining, maintenance, and manufacturing.

The field of MRSs is diverse, with research and development constantly progressing in a wide range of topics. We presented the experts with a list of topics, and asked them to choose the one they are most familiar with. An option to add their own topic was provided. Each expert was then asked to present the main challenges regarding the chosen topic. The two most frequently chosen topics are coordination/cooperation, control and planning (44.1%) and bio-inspiration, modular and swarm robots (25.0%). The most frequently mentioned challenges in coordination/cooperation, control and planning are related to system integration, particularly in unstructured environments; robustness of the system, both in terms of hardware and software/algorithms/behaviours; logistics, in terms of system deployment and energy; scalability, particularly in terms of how behaviours scale as more robots are added; and costs. Less frequently mentioned challenges include the reality gap; operation safety, in terms of humans, robots and tasks; and the lack of real-world testing. For the participants choosing the topic of bio-inspiration, modular and swarm robotics, the challenge most mentioned is the lack of accessibility to, and the high costs of, reliable and capable robot platforms. This is followed by logistics challenges, which are particularly severe when studying swarms of robots. One other challenge is the lack of realistic real-world applications, as many studies performed on this topic frequently fall back to simple and commonly used behaviours, such as aggregation or pattern formation, which would generally solve only a small part of a given real-world problem. The availability of capable, reliable and robust robot platforms is also the most frequently mentioned challenge on the topic of applications. For networked robots, communication, end-user acceptance and scalability are some of the challenges that were identified by the experts.

To be able to assess what can be achieved in the future, it is required to have information on what has been achieved in the past. With that in mind, survey participants were asked which area in MRSs they consider to have improved the most in the last 10 years. From the responses (89.7% of the participants replied), the areas most frequently mentioned are coordination/cooperation, control and hardware. From the responses to the previous question, we know that more than 40% of the participants are most familiar with coordination/cooperation, control and planning, and the challenges identified in these fields were mostly related to implementation issues. As a result, it is reasonable for coordination/cooperation, control and planning to be perceived as areas with significant development over the last 10 years. By contrast, hardware, in the form of available platforms, is one of the aspects frequently mentioned in the previous question (for multiple topics) that still presents a challenge to be overcome. Despite hardware being

one of the topics participants believe to have shown most development, further progress is vital, for example, to enable more comprehensive validation and/or deployment of solutions.

While current challenges and recent developments help determine possible future development paths, we also wish to identify promising areas for future research. Survey participants were asked what development specific to MRS they would like to see in the next 10 years. The most commonly mentioned developments were real-world testing (16.2%), real-world applications (13.2%) and coordination/cooperation (11.7%). The interest in real-world testing could be viewed as a consequence of an insufficient availability of hardware for implementation and testing. Although hardware has been identified as a topic of interest to develop in the next 10 years, the percentage of experts mentioning it is relatively small (4.4%). However, developments in this field are required to allow real-world testing to be performed. The interest in real-world applications could be viewed as a consequence of the change in perspective, as previously mentioned, where more theoretical work is being replaced with more practical work. Additionally, it is through the study of new applications that new challenges can be identified. Lastly, with new applications and more experimentation, the requirements for coordination/cooperation algorithms are expected to become more stringent, as MRS performance evaluations move from laboratory environments to realistic environments. Not only does the amount of recent development imply that coordination/cooperation is a growing topic, the development of new technologies, particularly in communication, is expected to enable more reliable coordination/cooperation of MRSs. However, significant work is required to reconcile new technologies and coordination/cooperation strategies.

One of the logistics concerns presented by the participants is the lack of adequately trained personnel. The participants were asked what they consider to be an effective activity for training/teaching staff or students to work with MRSs. The majority (85.3%) of the participants replied, resulting in the following proposed activities (ordered from most to least mentioned): hands-on experience; simulations; experimentation; courses; algorithm development; participation in competitions; study and discussions of recent literature; development of robotic systems; and summer schools.

In the last open-ended question, we asked participants to state the main difficulties faced in their current or past jobs when undertaking work related to MRSs. Different from what we expected, many of the responses concerned technical difficulties, with the most commonly mentioned ones being hardware reliability, logistics, complexity of experimentation, implementation issues, and communication. Non-technical difficulties that were reported include the costs and access to equipment and infrastructures, personnel training, space constraints, tool availability, and management of expectations in the field of study.

## 4 Conclusions

This paper reported a data-driven trend analysis of MRSs.

An exhaustive search of all papers published in 27 leading robotics journals from 2010 to 2020 revealed (i) a substantial growth in the number of papers on MRSs albeit the fraction of papers devoted to the topic decreased; (ii) for half of the journals, the fraction of papers devoted to MRSs exceeded 5%; (iii) six journals reported MRS research at TRL 6 or above; and (iv) a decline, in relative terms, in the use of laboratory demonstrations while theoretical work and real-world demonstrations remain in focus.

A survey of 68 leading experts revealed (i) that there is no consensus regarding what constitutes an MRS; (ii) the applications where MRSs are expected to be widely deployed in the short-, medium- and long-term; and (iii) the areas of development that require the most effort to move from single-robot systems to MRSs in real-world applications, including experimentation under realistic conditions, autonomy/resilience, system integration, trustworthiness and end-user acceptance. There was no consensus regarding how rapid the field will develop over the next decade, with some expected a faster development while others expected a slower development.

## Acknowledgments

This study was conducted as part of the Multi-Robot Systems project (Oct 2019–Dec 2020), which was funded by the Defence Science and Technology Laboratory (contract no DSTL-1000141223). The authors thank the 68 participants of the expert survey, and Alejandro R. Mosteo for feedback on the design of the expert survey.

## References

1. W. G. Walter, *The Living Brain*. WW Norton, 1953.
2. G. Dudek, M. R. M. Jenkin, E. Miliotis, and D. Wilkes, “A taxonomy for multi-agent robotics,” *Autonomous Robots*, vol. 3, no. 4, pp. 375–397, 1996.
3. A. Farinelli, L. Iocchi, and D. Nardi, “Multirobot systems: A classification focused on coordination,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 34, no. 5, pp. 2015–2028, 2004.
4. L. E. Parker, D. Rus, and G. S. Sukhatme, “Multiple mobile robot systems,” in *Springer Handbook of Robotics*. Springer, 2016, pp. 1335–1384.
5. F. Rossi, S. Bandyopadhyay, M. Wolf, and M. Pavone, “Review of multi-agent algorithms for collective behavior: A structural taxonomy,” *IFAC-PapersOnLine*, vol. 51, no. 12, pp. 112–117, 2018.
6. UK House of Commons, Technology and Innovation Centres, Science and Technology Committee, “Annex 1: Technology Readiness Levels,” 2010, <https://publications.parliament.uk/pa/cm201011/cmsselect/cmsctech/619/61913.htm#note221>.
7. American Society for Engineering Education, “Engineering and engineering technology by the numbers,” 2019, Washington, DC.