Safety for Mobile Robotic System: a Systematic Mapping Study from a Software Engineering Perspective

Darko Bozhinoski, Davide Di Ruscio, Ivano Malavolta ,
Patrizio Pelliccione, Ivica Crnkovic

PII: S0164-1212(19)30031-7
DOI: https://doi.org/10.1016/j.jss.2019.02.021
Reference: JSS 10302

To appear in: *The Journal of Systems & Software*

Received date: 18 September 2017
Revised date: 5 February 2019
Accepted date: 8 February 2019

Please cite this article as: Darko Bozhinoski, Davide Di Ruscio, Ivano Malavolta , Patrizio Pelliccione, Ivica Crnkovic, Safety for Mobile Robotic System: a Systematic Mapping Study from a Software Engineering Perspective, *The Journal of Systems & Software* (2019), doi: https://doi.org/10.1016/j.jss.2019.02.021

Highlights

- First study on the state of the art in safety for mobile robotic systems (MRSs)
- A reusable classification framework for methods and techniques for MRSs
- Classification of 58 studies w.r.t. trends, characteristics, and industrial adoption
- Discussion of future research challenges and implications on safety for MRSs

Safety for Mobile Robotic System: a Systematic Mapping Study from a Software Engineering Perspective

Darko Bozhinoski[a], Davide Di Ruscio[b], Ivano Malavolta[c], Patrizio Pelliccione[d,b], Ivica Crnkovic[d]

[a]IRIDIA, Université Libre de Bruxelles, Belgium
[b]University of L'Aquila, L'Aquila, Italy
[c]Vrije Universiteit Amsterdam, The Netherlands
[d]Chalmers University of Technology | University of Gothenburg, Department of Computer Science and Engineering, Sweden

---

Abstract

Robotic research is making huge progress. However, existing solutions are facing a number of challenges preventing them from being used in our everyday tasks: (i) robots operate in unknown environments, (ii) robots collaborate with each other and even with humans, and (iii) robots shall never injure people or create damages. Researchers are targeting those challenges from various perspectives, producing a fragmented research landscape.

We aim at providing a comprehensive and replicable picture of the state of the art from a software engineering perspective on existing solutions aiming at managing safety for mobile robotic systems. We apply the systematic mapping methodology on an initial set of 1,274 potentially relevant research papers, we selected 58 primary studies and analyzed them according to a systematically-defined classification framework.

This work contributes with (i) a *classification framework* for methods or techniques for managing safety when dealing with the software of mobile robotic systems (MSRs), (ii) a *map* of current software methods or techniques for software safety for MRSs, (iii) an elaboration on *emerging challenges and implications* for future research, and (iv) a *replication package* for independent replication and verification of this study. Our results confirm that generally existing solutions are not yet ready to be used in everyday life. There is the need of turn-key solutions ready to deal with all the challenges mentioned above.

Keywords: Software, Safety for mobile robots, Systematic mapping study

---

*Corresponding author. Email address: i.malavolta@vu.nl. Postal address: Vrije Universiteit, Faculty of Sciences, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands.

Email addresses: darko.bozhinoski@ulb.ac.be (Darko Bozhinoski), davide.diruscio@univaq.it (Davide Di Ruscio), i.malavolta@vu.nl (Ivano Malavolta), patrizio.pelliccione@gu.se (Patrizio Pelliccione), ivica@chalmers.se (Ivica Crnkovic)

1. Introduction

Robots are increasingly used in industry but also for tasks of our everyday life. In a recent book, Rise of the Robots [1], Martin Ford discusses the transition in robotics from special purpose robots, built to operate in highly controlled environments on a specific task, to general purpose robots that can operate in a heterogeneous environment, intermixed with humans, and perform a broad spectrum of tasks. Smart robots equipped with sensors and intelligent software promise to bring a new industrial revolution. According to Industrie 4.0 [2], we are in the middle of the 4th industrial revolution that is based on autonomous and smart Cyber Physical Systems (CPSs) [2], able to cooperate with each other and humans in a safe, autonomous, and reliable manner. The market for industrial robotics is expected to rise at a Compound Annual Growth Rate (CAGR) of 11,5% annually though 2021 and to reach $48.9 billion by 2021 [5]. The total smart robots market is expected to reach USD 7.85 billion by 2020, at an estimated CAGR of 19.22% between 2015 and 2020.

In this paper we focus on Mobile Robotic Systems (MRSs). This class of robots opens new long-term ambitions and business opportunities. Commercial drone revenue in Europe in 2017 was around $188 million, almost double the amount than in 2015 which was around $98 million [3]. Moreover, the total global Unmanned Aerial Vehicle (UAV) market is expected to grow from $ 20.71 billion in 2018 to reach $ 52.30 billion by 2025 [4, 5]. In a near future, there will be the need of customer-specific MRS solutions for a specific domain, such as: Homeland Security (e.g. coastal surveillance), Environmental Protection (e.g. emission monitoring and control), Protection of Critical Infrastructure (e.g. monitoring water and gas pipelines).

However, MRSs pose also important challenges: they need to be able to operate in uncontrollable and unknown environments, which are often shared with humans, and often they will be required to collaborate each other, and even with humans, to accomplish complex missions. Because of these challenges, these systems are both safety and mission critical. Safety criticality is an aspect of MRSs where failure or malfunction of the system may cause injure to people or severe damage to equipment/property, while mission criticality is another aspect of MRSs where a failure or malfunction may lead to an unacceptable loss of mission goals. Although robotic research has made huge progress in the last decades, the aforementioned functionalities and existing solutions seem to be not-yet-ready to be used in everyday life, and in uncontrollable and unknown environments often shared with humans [6], which will be shown as part of the conclusion of this study.

The **goal** of this study is to identify, classify, and evaluate the state of the art on safety for MRSs in terms of technical characteristics, potential for industrial adoption, and their challenges and implications for future research on safety for MRSs. The study exclusively focuses on software aspects.

In order to target our goal, we apply a well-established methodology from the medical and Software Engineering research communities called **systematic mapping** [7, 8]. The aim of a systematic mapping study is to provide an unbiased, objective and systematic approach to answer a set of research questions about the state of the art and research gaps on a given topic. A mapping study follows a well-defined and replicable principled process for both the search and selection of relevant studies, and the collected data and results synthesis tend to be more quantitative and qualitative [9, § 4.4]. Through our systematic mapping process, we selected 58 primary studies among 1,274 potentially relevant studies fitting at best three research questions we identified (see Section 3.1). Then, we defined a classification framework composed of more than 50 different parameters for comparing state-of-the-art approaches, and we applied it to the 58 selected studies. Finally, we analysed and discussed the obtained data for each parameter of the classification framework and how it fits in the research landscape about safety for MRSs.

Table 1: Main emerging challenges and implications for future research on safety for MRSs

Challenges	**Implications**
C1) **Single vs Multi-robots**: most of the studies surveyed in this paper focus on a single robot.	I1) There is the need of solutions addressing safety when multiple robots need to collaborate with each other in order to accomplish complex missions.
C2) **Openness and capability to cope with uncertainty**: many of the surveyed studies do not support adaptiveness capabilities and most of them are not able to deal with open systems, i.e., systems supporting the addition and removal of robots, human actors, etc. at runtime.	I2) The adoption of MRSs in tasks of everyday life would require more investigation in adaptiveness capabilities as well as in dealing with open systems.
C3) **Compliance to standards**: many domain-specific standards related to safety are currently available. Only a minority of approaches are compliant to standards that specifically target safety aspects.	I3) When developing a robotic system, specific standards have to be taken into account to make it compliant to them and safe for the considered application domain.
C4) **Rigor and Industrial Relevance**: the majority of evaluations in safety for robotic systems lack both rigor and relevance.	I4) New strategies are needed to ensure an adequate rigor and relevance when planning the evaluation of approaches for safety of robotic systems.
C5) **Research community on software engineering and robotics**: even though there is a growing interest, the community of software engineering for robotic systems is still not consolidated.	I5) The challenge for the research community is to promote a shift towards well-defined engineering approaches able to stimulate component supply-chains and significantly impact the robotics marketplace.

The **main contributions** of this study are:

- a reusable *comparison framework* for understanding, classifying, and comparing methods or techniques for safety for MRSs;
- a *systematic review* of current methods or techniques for safety for MRSs, useful for both researchers and practitioners;
- a discussion of *emerging research challenges and implications* for future research on safety for MRSs;

- a *replication package* containing detailed reports, raw data, and analysis scripts for enabling independent replication and verification of this study.

To the best of our knowledge, this paper presents the first systematic investigation into the state of the art on safety for MRSs. The results of this study provide a complete, comprehensive and replicable picture of the state of the art of research on safety for MRSs, helping researchers and practitioners in finding characteristics, limitations, and challenges of current research on safety for mobile robotic systems. The main emerging challenges and implications for future research on safety for MRSs are shown in Table 1.

Article outline. The article is organized as follows. In Section 2 we provide background notions for setting the context of our study by clarifying and discussing (i) mobile robotic systems, (ii) safety for mobile robotic systems, and (iii) existing studies on safety for MRSs. Section 3 describes in details the research methodology we followed for designing, conducting, and documenting the study. Data demographics is presented in Section 4, followed by a description of the obtained results in Sections 5, 6, and 7. We present limitations and threats to validity in Section 8. Related works are discussed in Section 9, whereas Section 10 closes the article with final remarks.

2. Background

2.1. Mobile Robotic Systems

Robots have been successfully deployed in industry to improve productivity and perform dangerous, tedious, or repetitive tasks [10]. In the literature, a variety of definitions exists defining the term "robot" [11, 12, 13]. All of them share the following concept: *a robot is an intelligent device with a certain degree of autonomy that contains sensors, control systems, manipulators, power supplies and software all working together to perform the required tasks.* Under this perspective, a **mobile robot** represents a robotic system consisting of a SW/HW platform carried around by locomotive elements and able to perform tasks in different contexts. The kind of locomotion that the robot is able to perform is primarily decided upon the environment (aquatic, aerial or terrestrial) in which the robot will be operating [14]. Mobility gives robots enhanced operative capabilities, but at the same time increases complexity and brings additional safety challenges.

In order to reduce the human involvement in scenarios that are characterized by repetitive and dangerous tasks (eg. natural catastrophes, nuclear power plant decommissioning, extra-planetary exploration, or less dangerous activities, such as delivery services, surveillance, and environmental monitoring), innovative technologies and approaches represented by mobile robotics are seen as particularly suitable for aiding in the process of replacement of the human beings with robotic systems. That will lead to a society where mobile robots will operate in a dynamic environment and perform the necessary tasks in these scenarios. But, if we want mobile robots to be widely accepted and adopted among the general public, it is fundamental to carefully consider safety aspects.

2.2. Safety for MRSs

One of the most important reasons for the success of industrial robotics is its assurance of a high degree of safety. However, industrial safety standards are focused on safety by isolating the robot away from people [15]. The new technological advancements in robotics enable robots to move from isolated environments to more unstructured and dynamic environments where they operate among people performing collaborative tasks beyond their explicitly preprogrammed behaviour. Hence, it is fundamental for safety aspects to be reconsidered and greatly enhanced at this point of time. We use the following definition of safety: *safety represents the absence of catastrophic consequences on the user(s) and the environment* [16]. In this context, *safety for MRSs* is defined as a property of the system that does not allow physical injury of people and loss or damaging to equipment/property in the environment. We consider as safety aspects all aspects of the system that involve prevention, removal, forecasting, and tolerance of faults and failures. Safety is a system property that should be addressed at every level of abstraction. In this study we focus on safety from software engineering perspective. It is difficult to distinguish between safety issues from different perspectives (e.g. software perspective in contrast to hardware, control theory or behavioural aspect) as it is difficult to draw a line between them. However, when safety is addressed from multiple aspects (e.g. software engineering, control theory, mechatronics), if the major contribution is towards software engineering principles and practises, we become inclusive and we are considering it in our study. This way we position our paper to help researchers in identifying design tools and methodologies for software for mobile robots that follow safety standards.

To address the increasing complexity and the needs of the variegated nuances of mobile robots, the robotics and automation industry are working towards the establishment of new international safety standards through the International Organization for Standardization (ISO) for robots and robot systems integration [15]. The current developed standards vary much as they depend on the particular application domains where the considered robotic systems are employed. The domain of personal care and agriculture is expanding rapidly. As a result, the **ISO13482** standard for *Safety requirements for personal care robots* and the **ISO18497** standard for *Safety of highly automated agricultural machines* have been developed. Another really important standard is ISO 15066, which focuses on the collaboration between people and robots. It specifies safety requirements for collaborative industrial robot systems and supplements the requirements and guidance on collaborative industrial robot operation. Commercialisation and adoption of mobile robots in dynamic environments will only occur if the safety aspects are considered and incorporated as first class elements in the design of the system. Establishing the guidelines and standards to regulate a safe use of these innovative technologies is the means to increase their trustworthiness and thereby their appreciation and use, not only in the research and business sectors, but also in the private social sphere. Certification bodies should assure safety certification that relies on a complete understanding of the system. However, for mobile robots that operate in dynamic environments it

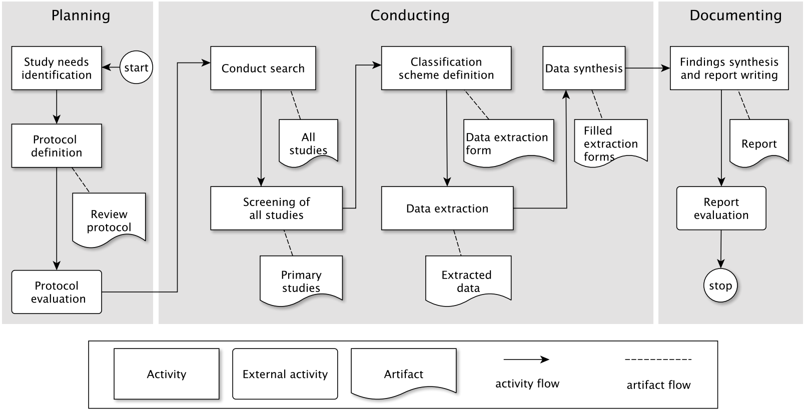

Figure 1: Overview of the whole mapping process

is quite challenging to consider all variants of the overall system due to their adaptive behaviour [17]. Recently, researchers have put their focus on the potential for using robots to aid humans outside strictly industrial environments, in more unstructured and dynamic ones [18]. The authors of [19] developed a safety module that integrates safety functions required for robots to work side by side with humans; it is compliant with international safety standards and Japanese law. It is strongly recommended to revise safety properties for MRSs in other application domains that will comply to identified international safety standards.

Finally, as of today we did not find any evidence that could help us in assessing the impact of existing research on *safety in mobile robots*. With this study we aim at helping researchers and practitioners in identifying the characteristics, challenges, and gaps of current research on this topic, its future potential, and its applicability in practice in the context of real-world robotic projects.

3. Study Design

Figure 1 shows the overview of the process we followed for carrying out this study. The overall process can be divided into three main phases, which are the classical ones for systematic mapping studies [8, 9]: planning, conducting, and documenting. In the following we will go through each phase of the process, highlighting its main activities and produced artifacts.

Planning. It is the first phase of our study and it aims at (i) establishing the need for performing a mapping study on safety for MRSs; indeed, as discussed also in Section 9, secondary studies exist on topics related to robotics safety like mechanical and controller design [20] and human-robot interaction [21, 22, 23], but none of them takes into consideration safety from a software engineering point of view; (ii) identifying the main research questions (see Section 3.1); and (iii) defining the review protocol detailing each step of the whole study. The output of the planning phase is a well-defined review protocol. In order to mitigate potential threats to validity, our review protocol has been circulated to external experts for independent review and we refined it according to their feedback[1].

Conducting. In this phase we carried out each step of the above mentioned review protocol. More specifically, we performed the following activities:

- *Conduct search*: in this activity we applied a search string to well-known academic search databases (see Section 3.2). The output of this activity is a comprehensive list of all the candidate studies resulting from the search.
- *Screening of all studies*: candidate entries has been filtered in order to obtain the final list of primary studies to be considered in later activities of the study. The basis for the selection of primary studies is the inclusion and exclusion criteria described in Section 3.2.
- *Classification framework definition*: we created a classification framework to compare the selected primary studies. The classification framework has been designed to collect data for answering the research questions of this study [9] and includes categories such as the level of abstraction in which safety is managed, compliance to standards, the scope and cardinality of hazards, etc. This activity will be described in more details in Section 3.3.
- *Data extraction*: in this activity we analysed each primary study, and we filled the data extraction form with the extracted information. Filled forms have been collected and aggregated in order to be ready to be analyzed during the next activity. More details about this activity will be presented in Section 3.4.
- *Data synthesis*: this activity focussed on a comprehensive summary and analysis of the data extracted in the previous activity. The main goal of this activity is to elaborate on the extracted data in order to address each research question of our research. The details about this activity are in Section 3.5.

Documenting. The main activities performed in this phase consist of (i) a thorough elaboration on the data extracted in the previous phase with the main aim of setting the obtained results in their context, (ii) the analysis of possible threats to validity, specially the ones identified during the definition of the review protocol (in this activity also new threats to validity may emerge), and (iii) the writing of a final report describing in details the design and results of this research.

3.1. Goal and Research Questions

We formulate the goal of this research by using the Goal-Question-Metric perspectives (i.e., purpose, issue, object, viewpoint [24]). Table 2 shows the

---

[1] We thank Richard Torkar (University of Gothenburh, Sweden) and Wasif Afzal (Mälardalen University, Västerås, Sweden) for their precious feedback on the review protocol.

result of the above mentioned formulation.

Table 2: Goal of this research

Purpose	Analyse
Issue	the characteristics and potential for industrial adoption
Object	of existing approaches for safety for MRSs
Viewpoint	from a researcher's and practitioner's point of view.

The goal presented above can be refined into the following main research questions.

- **RQ1**: *How do existing approaches address safety for MRSs?*

  Objective: to *identify* and *classify* existing approaches for safety in MRSs in order to build (i) a solid foundation for classifying existing (and future) research on safety for MRSs and (ii) an understanding of current research gaps in the field of safety for MRSs.

- **RQ2**: *What is the potential for industrial adoption of existing approaches for safety for MRSs?*

  Objective: to *assess* how and if the current state of the art on safety for MRSs is ready to be transferred and adopted in industry. Here we consider criteria such as the rigor and precision of the applied validation strategies (e.g., in-the-lab experiment, industrial application), the realism and scale of the performed evaluations, etc.

- **RQ3**: *What are the main emerging challenges for future research on safety for mobile robotics systems?*

  Objective: to *put into context* the results of RQ1 and RQ2 in order to *identify* the main challenges which will be faced by future researchers on safety for MRSs.

Answering those research questions will provide a solid foundation for understanding the state of the art on safety for MRSs, together with its research gaps and future challenges. The above listed research questions will drive the whole systematic review methodology, with a special influence on the primary studies search process, the data extraction process, and the data analysis process.

3.2. Search and Selection

The success of any systematic study is deeply rooted in the achievement of a good trade-off between (i) the coverage of existing research on the topic and (ii) having a manageable number of studies to be analysed [7, 8]. In order to achieve the above mentioned trade-off, our search strategy consists of two complementary methods: automatic search and snowballing. As shown in Figure 2,

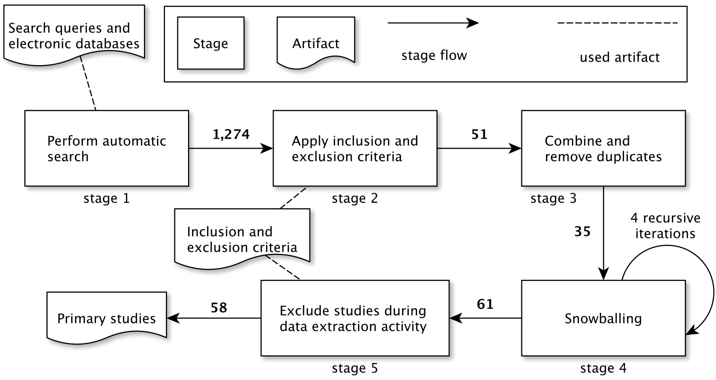

Figure 2: The search and selection process of this research

we designed our search strategy as a multi-stage process in order to have full control on the number and characteristics of the studies being either selected or excluded during the various stages. In the following we give a brief description of each stage of our search and selection process.

Stage 1. In this stage we performed automatic searches on electronic databases. In order to cover as much as possible relevant literature, four of the largest and most complete scientific databases were chosen as the sources of primary studies of this stage, namely: IEEE Xplore Digital Library, ACM Digital Library, SpringerLink, and ScienceDirect. The selection of these electronic databases is guided by (i) their high accessibility, (ii) their ability to export search results to well-defined, computation-amenable formats, and (iii) because they have been recognized as being an efficient means to conduct systematic literature reviews in software engineering [25, 26].

To create the search string, we break down our research questions into individual facets (population, intervention, comparison, outcomes, context - PICOC) as discussed in [27]. In our study, the PICOC elements that we identified are as follows:

- **Population:** mobile robotic systems;
- **Intervention:** approaches that address safety in mobile robotic systems;
- **Comparison:** not applicable;
- **Outcomes:** the classification framework populated with the identified primary studies;
- **Context:** academic peer-reviewed publications with a software engineering perspective.

Then we draw up a list of synonyms, abbreviations, and alternative spellings, which combined by logical ANDs and ORs gave the search string. Moreover, it is important to highlight that this study focuses on software aspects. This does not mean that safety in robotics is only a software aspect, but this is the focus of this study and the focus defines the boundary of the study itself. The obtained search string is given below and it has been tested by executing pilot searches on IEEE Xplore Digital Library.

> (mobile **OR** ground **OR** water **OR** fly* **OR** sail* **OR** unmanned **OR** self **OR** autonomous) **AND** (robot* **OR** vehicle*) **AND** (safe* **OR** fault **OR** failure) **AND** software

For the sake of consistency, the search strings has been applied to an identical set of metadata values (i.e., title, abstract and keywords) from all electronic databases. This stage resulted in a total number of 1,274 potentially relevant studies.

Stage 2. The main goal of this stage is to consider all the selected studies and filter them according to a set of well-defined inclusion and exclusion criteria. As suggested in [8], we decided the selection criteria of this study during its protocol definition, so to reduce the likelihood of bias. In the following we provide inclusion and exclusion criteria of our study. In this context, a study will be selected as a primary study if it will satisfy *all* inclusion criteria, and it will be discarded if it will met *any* exclusion criterion.

I1) Studies proposing an approach for safety for an MRS[2].

I2) Studies focussing on safety in MRSs from a software engineering perspective (e.g., no control theory or mechatronics studies, no studies focussing on hardware, etc.).

I3) Studies providing some kind of evaluation of the proposed methodology (e.g., via a case study, a survey, experiment, exploitation in industry, formal analysis, example usage).

I4) Studies subject to peer review [9] (e.g., journal papers, papers published as part of conference proceedings will be considered, whereas white papers will be discarded).

I5) Studies written in English language and available in full-text.

E1) Studies *exclusively* focussing on safety for industrial and other immobile robots.

E2) Secondary studies (e.g., systematic literature reviews, surveys) [9].

E3) Studies in the form of tutorial papers, short papers, poster papers, editorials, because they do not provide enough information.

---

[2] In the context of this research an *approach* can be considered as an organized set of methods and techniques, possibly supported by a tool [28].

In order to reduce bias, the selection criteria of this study have been decided during the review protocol definition (meaning that they have been checked by the two external reviewers).

In this stage, each potentially relevant study has been analysed in three phases. Firstly it has been analysed by considering its title, keywords, and abstract; secondly, if the analysis did not result in a clear decision, also its introduction and conclusions have been analysed; finally, we performed a comprehensive third manual step in which we read the full text of all considered studies (title, abstract, keywords, all sections and appendices, if any) in order to take the final decision about its inclusion in our set of primary studies. Two researchers have been involved during those phases and a third researcher has been involved in order to solve conflicts and take converge towards the final decisions, while avoid endless discussions [29].

In this stage, it is fundamental to select papers objectively. To this end, as suggested by [9], two researchers independently assessed a random sample of studies, then the inter-researcher agreement has been measured using the Cohen Kappa statistic; we obtained a Cohen Kappa statistic of 0.80, which is a good indication of the objectiveness of the performed selection. This stage resulted in a total number of 51 relevant studies.

Stage 3. In this stage all studies from the first stage have been combined together into a single set. Duplicated entries have been identified and merged by matching them by title, authors, year, and venue of publication. This stage resulted in a total number of 35 studies.

Stage 4. As recommended in guidelines for systematic studies, we extended the coverage of the previous stages by complementing the previously described automatic search with a snowballing activity. The main goal of this stage is to enlarge the set of relevant studies by considering each study selected in the previous stages, and focussing on those papers cited by it. More technically, we performed a *closed recursive backward and forward snowballing* activity [30]. From a practical point of view, we went through each selected study and we included also the relevant studies either cited by or citing it (based on Google Scholar [30]). The start set for the snowballing activity was composed of the 35 studies selected in stage 3. Then, we considered each paper in the start set and applied the same selection criteria discussed in stage 2 to each paper either cited by or citing it. If a paper was included, snowballing was applied iteratively until no new papers have been found. Duplicates were removed at each iteration of the snowballing activity.

This stage largely increased the number of potentially relevant studies, bringing it to 61. As a possible explanation of this fact, we noticed that researchers used a very heterogeneous terminology when writing the title, abstract, and keywords of their studies; this fact may negatively impact our automatic search, which may have missed some potentially relevant studies. We included the snowballing activity in order to mitigate this potential threat to validity. As a further confirmation, the study reported in [31] empirically observed that similar patterns and conclusions are identified when using automatic search and snowballing, especially when they are used in combination.

Stage 5. This stage has been performed in parallel with the data extraction activity. Basically, the idea is that when reading a study in details for extracting its information, researchers could recognize that it was out of scope, and so it has been excluded. This stage led us to the finalized set of 58 primary studies of our research, which is comprised of 58 entries.

3.3. Classification Framework Definition

One of the main contributions in our study is the classification framework, which consists of parameters that we identified as part of the protocol. We consider that these newly identified parameters can be reused in future studies to help authors of new methods and techniques to compare their contribution to existing ones. The different categories of our classification framework are described in more details in the following subsections. The **classification framework** is composed of three facets, each one dedicated to one of the RQ1 and RQ2 research questions (see Section 3.1). RQ3 does not have a dedicated facet in the classification framework since it is orthogonal to RQ1 and RQ2 and it aims at putting their results in the context of future emerging challenges on safety for MRSs. The classification framework also contains publication metadata (e.g., publication venues, authors, etc.), which have been collected for demographics purposes (see Section 4).

3.3.1. How safety for MRSs is managed (RQ1)

Since research question RQ1 is at the core of our research, the creation of its corresponding facet in the classification framework demands a detailed analysis of the contents of each primary study. In light of this, we followed a systematic process called *keywording* [32] for building this facet of our classification framework. Keywording aims at reducing the time needed in developing a classification framework and ensures that it takes the considered studies into account [32].

As shown in Figure 3, keywording is done in two steps:

1. *Collect keywords and concepts*: we collected keywords and concepts by reading the abstract of each primary study. When all primary studies have been analysed, all keywords and concepts have been combined together to clearly identify the context, nature, and contribution of the approach. As suggested in [32], when the abstract of a primary study was not informative enough, then we analysed also its introduction and conclusion sections. Considering that the authors of the primary studies may use different terms for same concepts and same terms for different concepts, in this phase we kept all keywords and concepts to ensure consistency and compatibility. The output of this stage is the set of keywords as they have been used in each primary study.

2. *Cluster keywords and form categories*: when keywords and concepts have been collected, then we performed a clustering operation on them in order to have a set of representative clusters of keywords. We identified

the clusters by applying the open card sorting technique [33] to categorize keywords into relevant groups. More specifically, we considered all the keywords and concepts collected in the previous phase and iteratively grouped them together until a saturation of all the concepts has been achieved and all primary studies were analyzed. In order to minimize bias, this step has been performed by two researchers and the results have been double-checked by the other two researchers. The output of this stage is the classification framework containing all the identified clusters, each of them representing a specific aspect of safety for MRSs. The specific categories emerging from the keywording process are described in Section 5.

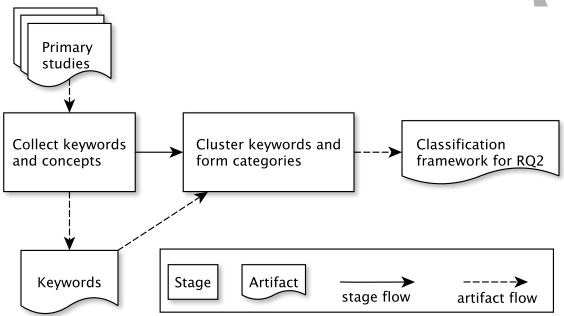

Figure 3: Overview of the keywording process

Moreover, we collected also data related to the *main research contribution* and *application field independence* of each primary study. The categories for research contributions are derived from [32] and include values such as “method”, “architecture”, “tool”; they are discussed in details in Section 5.1. For what concerns application field independence, while piloting this study we noticed that in the discussion of related work of some papers authors were referring to both domain-specific approaches and generic ones; based on this, we decided to categorize our primary studies about whether they are independent with respect to any application field (e.g., abstract approaches orthogonal to any application field) or not (e.g., approaches that are specifically tailored to self-driving cars, agriculture, environmental monitoring).

Since this research question is of key importance for this survey, we made a pre-study in order to classify existing works on safety mechanisms. The pre-study consists in analysing three recent surveys on MRS safety from 2017, namely [34, 35, 36] and we extracted the parameters they have used in their classification schema and we used on our primary studies. For each of the primary studies, we collected in a spreadsheet a record for each parameter. Each cell in the record represents a boolean value that give information if the primary study is addressing a particular aspect represented by the parameter extracted from the surveys.

All three surveys are secondary studies that address MRS safety from different domain, having different perspective and conclusions. [36] focuses on classification schema for methods for safe human-robot interaction, [34] is a survey on dependability techniques used for increasing safety in MRS addressing large scope of application domains and [35] reviews and evaluates model-based algorithms for real-time collision detection, isolation, and identification focusing on control strategies for safe robot reaction. As we see all the surveys address safety from a different perspective. We extracted all the parameters they have used in all three surveys and we used this classification schema on our primary studies. For each of the primary studies, we collected in a spreadsheet a record for each parameter. Each cell in the record represents a boolean value that gives information if the primary study is addressing a particular aspect represented by the parameter extracted from the surveys. All parameters have been described in Table 3.

3.3.2. Potential for industrial adoption (RQ2)

To answer this research question we performed an analysis of qualitative data. To perform the analysis we used the already presented keywording method, and then we analysed and summarized the potentials for industrial adoption that have been highlighted in the papers. The parameters that we considered are:

- *applied research method*: here we distinguished between approaches validated in a controlled setting (or in the lab) and approaches evaluated in real-world (industrial) contexts;
- *validation/evaluation strategies*: here we extracted the strategies applied for assessing the proposed approaches (e.g., real deployment, simulation-based, proof of concept), independently of whether they are performed in the context of validation or evaluation research;
- *technology readiness level (TRL)*: it has been proposed by the Horizon 2020 European Commission for the 2014/2015 work program[3], the TRL is a metric for measuring the maturity of a given technology;
- *rigor and industrial relevance*: we measured the precision, exactness and realism of the evaluation of each primary study by applying the rigor and industrial relevance metrics proposed by Ivarsson and Gorschek [37];
- *industry involvement*: whether each primary study has been carried out only by academics, practitioners (or a mix therof) for understanding how researchers and practitioners collaborate on safety for MRSs.

---

[3] http://ec.europa.eu/research/participants/data/ref/h2020/wp/2014_2015/annexes/h2020-wp1415-annex-g-trl_en.pdf

Table 3: Classification parameters proposed by other secondary studies

Survey	Parameters	Description
A Survey of Methods for Safe Human-Robot Interaction [36]	Reactive Safety	if it is reactive (not performing any planning)
	Proactive Safety	if it is proactive (producing plans to address specific safety-related issues)
	Proactive Safety with prediction	if it can anticipate the actions and movements of the rest of the team of mobile robots or people
	Psychological safety	if it takes consideration of psychological factors
Safety-critical advanced robots: A survey [34]	Fault prevention	if it prevents the occurrence or introduction of faults, including techniques coming from system engineering and good practices from system designing
	Fault removal	if it reduces the number and severity of faults
	Fault forecasting	if it estimates the present number, the future incidence, and the likely consequences of faults.
	Fault tolerance	if it avoids service failures in the presence of faults using redundancy, error detections
Robot Collisions: A Survey on Detection,Isolation, and Identification [35]	Precollision	if it discusses collision avoidance strategy
	Detection	if it has ability to understand if a system collision occurred
	Isolation	if it understands the impact of the collision
	Identification	if it understands the impact of the collision
	Classification	if it has capability to understand the nature of the collision
	Reaction	if it provides strategies for the system to react purposefully to a collision event
	Post-collision	if it discusses strategies how the robot will proceed after a safe state has been reached

3.3.3. Emerging challenges for Future Research (RQ3)

To answer this research question we followed a similar strategy to the one used for RQ2. We basically analyzed all the primary papers with the aim of collecting all the challenges that have been highlighted in such papers, and then we summarized the results that emerged.

3.4. Data extraction

As already said, the classification framework is the base of the data extraction form, i.e., a well-structured form to store the data extracted from each primary study. For each of these studies, we collected in a spreadsheet a record with the extracted information for subsequent analysis. As suggested in [9], the data extraction form (and thus also the classification framework) has been independently piloted on a sample of primary studies by two researchers, and iteratively refined accordingly. Once the data extraction form was setup, we considered each primary study and its corresponding data extraction form has been filled with the extracted data.

In order to validate our data extraction strategy, 10 primary studies have been randomly selected and two researchers checked whether the results were consistent, independently from the researcher performing the extraction. In this context, each disagreement has been discussed and resolved, with the intervention of a third researcher, when necessary.

3.5. Data Synthesis

This activity involved collating and summarising the data extracted from the primary studies[8, § 6.5] with the main goal of producing the actual map of current research on safety for MRSs. When possible, in this research we applied both quantitative and qualitative analysis methods, depending on the nature of each specific parameter of the classification framework.

For each parameter of the classification framework we divided our *quantitative* analysis on two main steps: (i) we counted the number of primary studies falling in relevant categories in the context of the specific parameter and (ii) we aggregated and visualized the extracted information to better clarify similarities and differences between the primary studies.

For what concerns the analysis of *qualitative* data, we used the already presented keywording method for identifying also the possible values of each parameter of the classification framework, and then we analysed and summarized the trends and collected information in a quantitative manner.

Finally, we carried out a narrative synthesis of the results obtained both quantitatively and qualitatively; this step allowed us to (i) perform an evidence-based interpretation of the main findings coming from the previous analyses and (ii) extract the main challenges and implications for future research. Narrative synthesis refers to a commonly used method to synthesize research in the context of systematic reviews where a textual narrative summary (i.e., by using words and text) is adopted to explain the characteristics of primary studies [38], alongside or instead of a statistical analysis [39, 40]. In the context of our study, for each parameter of our classification framework we firstly summarized it from a quantitative perspective (i.e., statistical summary) and then we complemented such quantitative analysis by applying the general framework for narrative synthesis proposed in [38], namely: (i) we developed a theory about the specific values of the parameter by tabulating the results and iteratively performing content analysis sessions, (ii) we developed a preliminary synthesis of findings based on the quantitative analysis, (iii) we explored potential relationships in the data (i.e., horizontal analysis), (iv) we assessed the robustness of the synthesis by critically reflecting on the synthesis process and checking the obtained synthesis with authors of primary studies.

4. Demographics

This research considers a set of 58 primary studies, each of them published in different years and venues. Figure 4 shows the distribution of the primary studies over the years and by the type of venue where they have been published[4]. The obtained data clearly shows a growing trend in terms of **publication intensity**, with most of the studies published in the very recent years; specifically, 46 studies over 58 have been published from 2009 to 2016 (with an average of more than 5 publications per year), where 17 studies have been published only in 2015 and 2016. If we look at the publication numbers before 2009 we have a drop to less than one publication per year. These results are a confirmation of the growing scientific interest on safety for mobile robotic systems, specially in the last years. The motivations behind such a publication trend can be manifold including the growing interest about autonomous vehicles[5] and the increasing funding opportunities for developing robotic systems to be employed both in industrial and in domestic contexts[6].

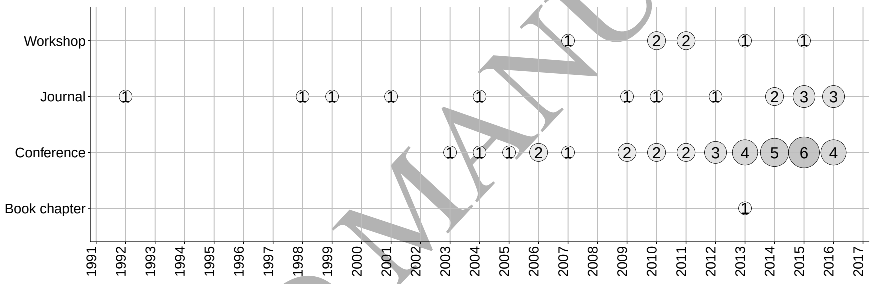

Figure 4: Distribution of primary studies over the years - results

More on a historical perspective, the first study on safety for mobile robotic systems (P11) has been published in the Applied Intelligence international journal in 1992. In P11 the authors proposed an automated diagnostic method for keeping an autonomous underwater vehicle operational for several weeks without human intervention. The approach was based on a distributed fault-tolerant control system aiming at managing unpredicted faults by preserving its overall performance level. The approach makes the assumption that the normal behaviour of each component is available at design time.

We also classified the primary studies by (i) type of publication and (ii) targeted publication venues. As shown in Figure 4, the most common **publication type** is conference paper (34/58), followed by journal papers (16/58), workshop

---

[4] Our search activity covers the research studies published until January 2017, thus we potentially have only partial data for 2016.

[5] https://www.gartner.com/smarterwithgartner/the-road-to-connected-autonomous-cars/

[6] https://www.gartner.com/doc/3418843/market-trends-personal-assistant-robots

papers (7/58), and finally book chapters (1/58).

Table 4: Targeted publication venues

Venue Acronym	**#Studies**	**Studies**
Intelligent Robots and Systems (IROS)	6	P1, P5, P20, P22, P32, P58
International Conference on Robotics and Automation (ICRA)	4	P6, P36, P40, P54
International conference on Automated Software Engineering (ASE)	3	P14, P21, P55
IEEE Transactions on Robotics and Automation (TRA)	2	P2, P7
Robotics and Autonomous Systems (Journal)	2	P10, P41
International Journal of Robotics Research (IJRR)	2	P19, P35
Conference Towards Autonomous Robotic Systems (TAROS)	2	P43, P45
IEEE Conference on Emerging Technologies and Factory Automation (ETFA)	2	P18, P56
IEEE International Symposium on High-Assurance Systems Engineering (HASE)	2	P21, P55
International Conference on Advanced Robotics (ICAR)	2	P24, P37
International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR)	2	P12, P27
Others	32	P28, P11, P26, P17, P8, P46, P31, P25, P47, P50, P9, P29, P49, P3, P34, P44, P39, P33, P53, P23, P16, P4, P13, P30, P57, P48, P15, P55, P38, P51, P52, P42

In Table 4 we report the **publication venues** that hosted more than two publications (the last row of the table is an aggregate of all the publication venues with two or less publications). What strikes the eye is the extreme fragmentation of the targeted publication venues (43 unique venues for 58 publications). Nevertheless, we can observe that the most targeted venues (i.e., the ones targeted by at least two primary studies, see Table 4) are quite homogenous and dedicated to robotics, autonomous systems, automation, and high-assurance systems. It is important to note that with Table 4 we are not aiming at establishing which publication venue is the most related to safety for MRSs; indeed, the size and frequency of conferences and journals may influence the numbers reported in the table (e.g., a yearly conference has potentially more safety-related publications w.r.t. a biannual conference). Nevertheless, given their focus on aspects related to safety for MRSs, we can consider the venues reported in Table 4 as good candidates for future publications on this area.

In the following we present the results of this study for answering our research questions (see Section 3.1). For each parameter of our classification framework we report both quantitative data and an interpretation of the obtained results.

5. How safety is managed (RQ1)

This section aims at identifying and classifying existing methodologies that address safety in mobile robotic systems.

In Section 3.3.1 we explained that in order to provide a classification framework we performed keywording that produces as output the formation of categories of the classification framework. Keyworking is a standard technique and more information might be found in Section 3.3.1 and in [32]. Roughly speaking, we collected all keywords across all papers and we group them together into meaningful groups. The resulting groups are then clustered into attributes and values (with different possible levels of hierarchy). The data extraction form is available at [41]. Figure 5 shows a graphical and tree-based representation of the categories in the classification framework. It is important to highlight that the categories that have been identified for safety management from the analysis of our primary studies through keywording and by following the process described in Section 3.3.1. What emerges from this classification is that, for designing a solution for safety management we need to consider also other aspects, like the nature of hazards, the characteristics of the system, whether models are used or not, and the involved standards, if any.

According to the classification framework and the summary of the categories in Figure 5, research question RQ1 has been decomposed into more detailed subquestions. Therefore, we discuss about:

- *safety management*: how the proposed approach considers safety-related aspects (e.g., specific mechanisms for safety, the level of abstraction, whether safety is treated as first class element of the approach or not, etc.) as shown in Figure 6;
- *system characteristics*: the features of the systems supported by the proposed approach (e.g., cooperative versus local adaptation, the type of robots, their cardinality, etc.);
- *models:* it is about the models[7] of the system and their features (i.e., whether the proposed approach is based on model-based techniques, the purposes of the used models);
- *standards*: the standards to which the proposed approach is compliant (e.g., IEC61508, ISO10218);

---

[7] It is important to remark that in this paper, with the term model we refer to specifications defining the different software aspects of the system being developed (e.g., requirement, component, and deployment specifications). Thus we do not refer to other kinds of models like 3D, mathematical, and physical ones that are considered by the robotic community.

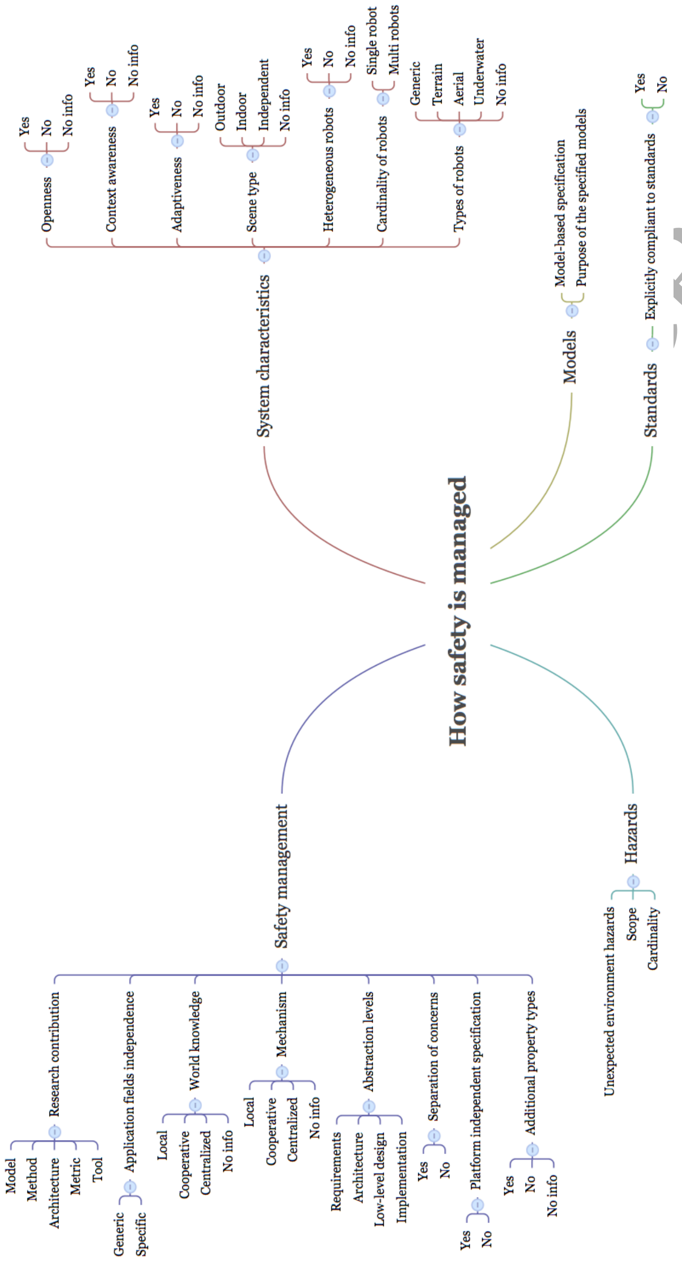

Figure 5: How safety is managed

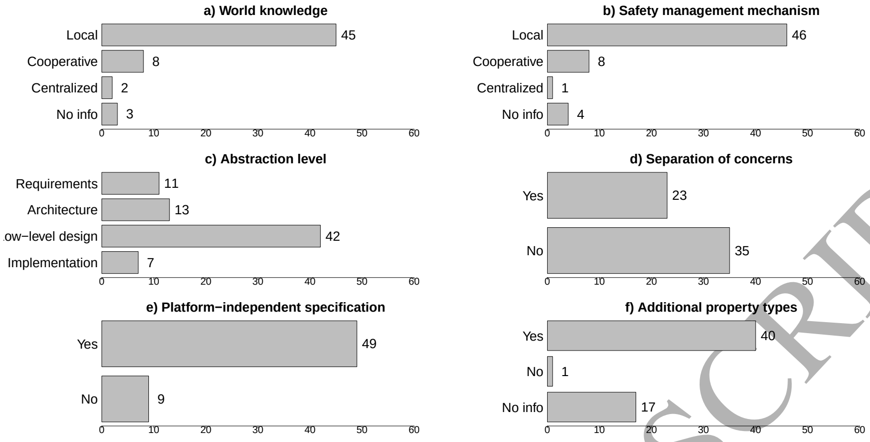

Figure 6: Safety management - results

- *hazards*: about the characteristics of the hazards managed by the approach (i.e., whether they are unexpected, their scope and cardinality).

In addition to that, by following what discussed in Section 3.3.1, in the highlights of RQ1 (end of this section) we classified the primary studies with respect to parameters of other secondary studies that we discovered in a pre-study, as described in Section 3.3.1.

5.1. Safety management - Research contributions

In order to characterize where researchers are focussing their efforts, we extracted the main research contribution of each primary study. Categories of research contributions are derived from [32], and can be one or more of the alternatives shown in Table 5.

The results of our analysis are shown in Figure 7.a. It does not come as a surprise that the main contribution of the majority of primary studies is a *method* to address specific concerns about safety for MRSs (43/58); this result does not come as a surprise since our inclusion criterion I1 is explicitly dealing with studies proposing either a method or a technique for safety. The second most recurrent research contribution is *architecture* (21/58); those studies present the fundamental concepts or properties related to the safety of an MRS by reasoning on its elements, relationships, and in the principles of its design and evolution [42]. This result is interesting since it confirms that safety has been treated as a system-level property by researchers, and that considering safety at a higher level of abstraction is a valuable and effective strategy for attacking the problem. Other studies contribute with the information, representations, and abstractions for safety of MRSs (*model*, 11/58), and developed tools or prototypes for safety of MRS (*tools*, 9/58). As a final consideration, no primary study has as main contribution *metrics*, indexes, or measures to assess

Table 5: Types of research contribution (adapted from [32])

Research contribution	**Description**
Model	Presents information, representations, and abstractions to be used in safety for MRSs.
Method	Presents general concepts and working procedures to address specific concerns about safety for MRSs.
Architecture	Presents the fundamental concepts or properties of an MRS embodied in its elements, relationships, and in the principles of its design and evolution [42].
Metric	Presents specific indexes and measures to assess certain properties of safety for MRSs.
Tool	Presents any kind of developed tool or prototype related to safety for MRSs.

certain properties of safety of MRSs. By following old adage that *what gets measured gets managed*, working on safety-specific metrics for MRSs can be an 625 added value for the field and surely an interesting research gap to be filled by future research.

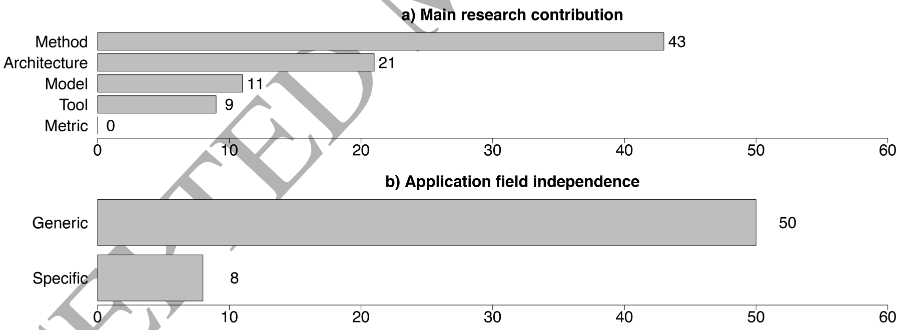

Figure 7: Types of research contribution (a) and application field independence (b) - results

5.2. Safety management - Application fields independence

As shown in Figure 7.b, almost all the primary studies are *generic* with respect to any application field. This means that those studies are kind of orthogonal and can be applied to some extent to different types of robots, tasks 630 to be performed, operational contexts, etc. For example, the authors of P9 achieved generality by applying the well-known abstraction and automation principles of the Model-Driven Engineering paradigm (MDE, [43]). By quoting their own words, their approach *directly enables an implementation-independent reuse of the safety-related part of a robot controller between different releases, since the RuBaSS declaration does not need to change when the underlying software changes (except that names shared between RuBaSS rules and component interfaces must be kept consistent). Moreover, the infrastructure can be reused in a range of products: the code generator can be directly reused whereas low-level interfaces to sensors and actuators will be specific to each robot. Safety-related customisation for the products is thus mainly achieved at the higher level, using the safety language* (P9).

Application-specific approaches have been proposed in 8 primary studies (namely, P7, P9, P11, P39, P42, P50, P54, and P56), with application fields ranging from health to domestic or industrial robotics.

It is important to know that application field independence is strongly related to the level of abstraction of a given approach. Specifically, a higher level of abstraction can result in a higher potential for reuse across domains, thanks to the abstraction from the low-level details and constraints of a specific domain. Also, if an approach is independent from a specific domain, then potentially it may be used by a wider community, leading to higher potential for cross-fertilization across disciplines (e.g., an obstacle avoidance algorithm for planetary exploration may be used and adapted for terrestrial exploration), or even more bugs discovered (and potentially fixed) in the tool supporting the approach. Nevertheless, having an approach specifically tailored to a given domain (e.g., exploratory robots in wild areas) allows engineers to be more specialized when solving domain-specific issues (e.g., how to manage the interaction with wild animals), potentially raising the chances of industrial adoption in the short term.

5.3. Safety management - World knowledge

It is important to identify the knowledge of the robot of the environment in which the robot will operate. When we deal with multi-robots, the various robots might share the knowledge about the environment in different ways. We believe these are important aspects that should be taken into account for having robots able to perform everyday tasks in environments that, increasingly, will be uncontrollable and only partially known.

As shown in Figure 6.a, most of the approaches (45/58) rely on a local knowledge of the environment. This means that the knowledge about the environment (including other robots involved in the mission) is local to each robot, without mechanisms to share knowledge between different robots. 2 approaches have a centralized world knowledge, meaning that the knowledge of the overall system is maintained by a centralized entity. 8 approaches have cooperative world knowledge and this means that there are mechanisms to share knowledge between different robots that take part in the mission.

It is important to note that only two approaches with local knowledge involve multi-robots, namely P43 and P51. This explains why we have a majority of approaches that rely on local knowledge. In general, we might say that having a centralized world knowledge in multi-robot systems might hamper the adoption of decentralized algorithms for (re)planning, issues resolution, and so on.

Managing the uncertainty of the environment where the considered robot has to operate is an orthogonal aspect, which is cross-cutting to those previously mentioned. Even though having the availability of a complete model of the environment represents the ideal situation, in practice only partial and limited world models are possibly available and consequently, specialized techniques are needed to permit robots to work with uncertain world knowledge. For instance, in [44] authors propose an approach for modeling cooperative intelligent vehicles by means of modeling constructs enabling the specification of uncertainty degrees for attributes of the modeled objects. In [45] authors propose an approach to support world modeling for autonomous systems. The main characteristic of the proposed technique is that *"it models uncertainties by probabilities, which are handled by a Bayesian framework including instantiation, deletion and update procedures"*. Recently, a novel approach has been proposed to deal with uncertainty of software models, by focusing on measurement uncertainty, and confidence [46]. However, dealing with uncertainty is a very challenging problem and an in-depth treatment of it is beyond the scope of this section, which is more focused on the way world knowledge is managed (e.g., locally or in a cooperative manner) and not on its content.

5.4. Safety management - Mechanism

Concerning this parameter we do not list the different mechanisms, but we categorize them as local, centralized or cooperative. A mechanism is local if it is conceived to work on single robots, without any cooperation, centralized if there is an entity managing the safety aspect of the system, or cooperative if safety mechanisms involve a cooperation between different robots. As shown in Figure 6.b, most of the approaches (46/58) adopt local safety mechanisms, i.e. safety mechanisms that are conceived to work on single robots, without any cooperation. This is expected since, as highlighted in Section 5.3, most of these approaches focus on single robots. The exceptions are P43 and P51 that deal with multiple robots even though they have local safety mechanisms, and P54 that has both local and centralized safety mechanisms. As can be seen in the figure only 1 approach has a centralized safety management mechanism. Instead, 8 approaches rely on cooperative safety mechanisms, meaning that safety mechanisms involve a cooperation between different robots. Finally, 4 approaches provide no information about this aspect.

5.5. Safety management - Abstraction levels

When developing complex systems, abstraction is a key concept to master complexity. In software engineering, the systems to be developed are analyzed at different levels of complexity by focusing on a few issues and aspects at a time. As shown in Figure 6.c, the abstraction level of the safety management spans from requirement till implementation. A requirements value means that safety is considered when eliciting/specifying the requirements of the system (e.g., generic safety rules written in a non-technical way). Architecture means that safety is considered at the architectural level (e.g., they talk about architectural tactics, styles, architectural patterns, system infrastructure, communication topology, etc.). Low-level design means that safety is considered at the design level (e.g., design patterns, design models, etc.). Finally, implementation means that safety is considered at the source code, programming level.

The majority of the approaches works at the design level that seems to be the most appropriate level to reason and deal with safety management. The design level is followed by the architecture level. In fact, as shown in Table 6, 7 approaches address safety at the implementation level and among them only 3 approaches exclusively address safety at the implementation level, 3 approaches address safety also at the design level, and 1 at the architecture level. This testifies that it might be difficult to manage safety directly at the implementation level and it is more profitable to deal with it at more abstract levels.

Table 6: Safety management - Abstraction levels

Level(s)	**Number of studies**
Requirements	3
Requirements + Low-level design	8
Architecture	9
Architecture + Low-level design	3
Architecture + Implementation	1
Low-level design	28
Low-level design + Implementation	3
Implementation	3

5.6. Safety management - Separation of concerns

As shown in Figure 6.d, for the majority of the approaches (35/58), the management of safety-specific issues (e.g., safety rules) is not kept separated from the functional management of the robots (e.g., the mission). Keeping a separation of concerns means for instance that the approach prescribes a special layer for managing safety, which is totally separated from the rest of the system. Managing complex missions requires a clear separation of concerns between safety and other aspects of the system. We consider that safety-specific objectives should be separated from the rest of the system because the nature of the safety objectives is different to the other objectives (e.g., mission objectives). Safety is considered as a first class concern in MRSs which means that MRS should always satisfy the safety objectives, while the other concerns (e.g. mission concerns) can be partially satisfied. That way a safety engineer can focus on definition of safety-specific mechanisms that are generic and independent from the functional behaviour of the system, while, for example, an operator can focus on the mission functional specification.

5.7. Safety management - Platform independent specification

As shown in Figure 6.e, for the high majority of the approaches (49/58), the specification of safety-specific aspects (e.g., safety constraints, properties, rules, invariants specifying assumptions about hardware) is independent from the underlying platform (e.g., ROS, hardware, operating system, etc.). This is a good characteristic of the platform since this can enable reusability of software modules across various platforms.

5.8. Safety management - Additional property types

Table 7: Additional property types (as reported in the primary studies)

Property	**#Studies**	**Studies**
Performance	13	P1, P2, P4, P10, P11, P13, P18, P23, P27, P28, P31, P34, P58
Functional correctness	12	P17, P41, P48, P49, P50, P51, P52, P53, P54, P55, P56, P57
Reliability	7	P18, P30, P31, P37, P38, P40, P46
Dependability	5	P10, P14, P20, P24, P31
Usability	5	P16, P21, P22, P27, P32
Robustness	4	P24, P35, P36, P37
Availability	3	P10, P22, P35
Effectiveness	3	P1, P11, P35
Reusability	3	P16, P27, P32
Efficiency	3	P1, P2, P10
Modularity	2	P16, P27
No additional property type	1	P3
Integrability	1	P21
Validity	1	P21
Applicability	1	P21
Maintainability	1	P45
Complexity	1	P45
Flexibility	1	P45
Expressiveness	1	P45
Upgradeability	1	P16
Reusability	1	P27
Repeatability	1	P32
Security	1	P22

As shown in Figure 6.f, most of the approaches deal with properties that are different from safety. In fact, 40 approaches deal with additional properties, only one approach is exclusively focused on safety, P3, and 17 papers do not provide information. Table 7 shows the additional properties and adds a reference to primary studies that are addressing the specific properties. There is a big variety of additional properties that are addressed by the primary studies - 22 different additional properties considered by the 40 primary studies that consider additional properties. Performance is the most addressed property, followed by functional correctness. The motivations behind the interest on performance when managing safety of MRSs can be manifold, including the need of improving the non-functional properties of the software and hardware components that are involved when reacting to unexpected events. Similarly, functional correctness is an additional property to be addressed for example when developing monitors that can detect conditions that may lead to failures and thus need to take corrective actions.

5.9. System characteristics - Openness

In the context of this study, by open systems we mean those systems that allow for entrance and exit of entities during mission execution [47]. Openness can improve the dynamicity of the MRS, for example by allowing to let new robots with better or new functionalities (or new human actors) to get into the MSR or to let robots that have completed their tasks to exit the MSR. As shown in Figure 8.a, most of the approaches are unable to deal with open systems (only 5 approaches, namely P2, P22, P48, P49, P53, are able to deal with open systems). This implies that most of the approaches that have been proposed are not able to manage safety once the system evolves in terms of addition or removal of robots and/or other types of agents, including humans. This is indeed an interesting research direction since systems of the near future will be necessarily characterised by openness, and it is often impossible to assess at design time the exact boundaries and topology of the system.

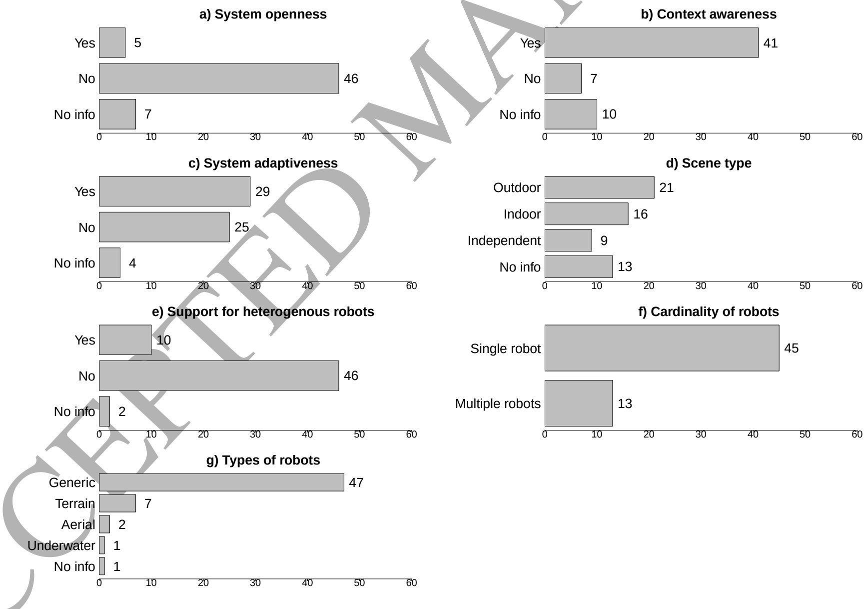

Figure 8: System characteristics - results

5.10. System characteristics - Context awareness

As can be seen in Figure 8.b, most of the approaches (41/58) deal with systems (including the robots) that are able to understand some key properties about the operational context of the robots (e.g., presence of obstacles, existence of other robots, etc.). 10 out of 58 approaches do not provide information. Context awareness is another important characteristic to enable the adoption of robots in real life scenarios, where often the operational environment is (partially) unknown and uncontrollable.

5.11. System characteristics - Adaptiveness

Figure 8.c shows that 29/58 approaches have adaptiveness capabilities. In the context of this study, *adaptiveness* means that the system (including the robots) is able to adapt (e.g., behaviour adaptation, trajectory recalculation, goal renegotiation) in order to find a solution depending on some change in the context of the mission being performed (e.g., unexpected obstacles, software/hardware failures, mission redefinition by a human actor). If all the possible adaptation alternatives are defined a-priori and analysed, then the system at runtime should be simply able to "safely" switch from one alternative to another. If at runtime the system will encounter unplanned situations, then there should be the transition towards emergency behaviours, opportunely planned and analysed. Adaptiveness might also require the use of learning techniques that, instead of switching among pre-defined alternative behaviours, will calculate at runtime what to do, for instance, by using machine learning algorithms. These techniques are very promising for dealing with uncertainty and partial knowledge in the environment, however, the use of machine learning for safety critical systems is still open [48]. 25/58 approaches do not support this functionality, and 4 approaches provide no information. Adaptiveness might be considered in conjunction with context awareness since awareness of the context is a required capability in order to support adaptiveness.

5.12. System characteristics - Scene type

This parameter aims to show how much of the safety approaches are tailored for specific scene types and how much of them are independent from the type of scene where the MRS is performing its mission. Figure 8.d describes the ability of the system to work indoor (21/58), outdoor (16/58), or independent of the scene (9/58). Some approaches provide no information in this concern (13/58). Please notice that we categorised an approach as independent only if the approach explicitly mentions about its independence ability. In conclusion, the majority of the safety approaches are tailored to systems that perform in a specific scene type (indoor or outdoor) instead of having a more generalized safety approach.

5.13. System characteristics - Heterogeneous robots

Another peculiar system characteristic is the capability of managing teams consisting of robots of different types (e.g., robots for grabbing objects, for video streaming, sensing and discovering relevant information). According to Figure 8.e most of the analyzed systems (46/58) do not have the capability of managing heterogeneous robots. Only 10 systems provide users with such a functionality, whereas 2 analyzed systems do not provide a clear statement about that. Hence, most safety approaches that are addressing team of robots are focused on homogeneous robots.

5.14. System characteristics - Cardinality of robots

Missions can be executed by one or more robots. Indeed the management of different robots introduce additional challenges mainly related to their collaboration and coordination. As shown in Figure 8.f most of the analyzed systems (45/58) support missions performed by a single robot (e.g., self-driving car), while few of them deal with the management of multiple robots. Hence, main focus on safety approaches have been single robots. Researchers should consider proposing solutions that will address safety on a team level.

5.15. System characteristics - Type of robots

This parameter can have values in the set {TERRAIN, UNDERWATER, AERIAL, ACQUATIC, GENERIC}. If the authors of a primary study explicitly claim that their proposed approach is specific to a type of robots (e.g., UAVs), then we set the value of this parameter to the family of the specific type of robot (e.g., AERIAL); if the authors of a primary study claim that their proposed approach is independent of the type of robots, the value of this parameter has been set to GENERIC. In order to manage different kinds of missions it is preferable that the used system provides users with functionalities that are robot independent. According to the performed analysis, 7 out of 58 analysed systems are specific to terrain robots (see Figure 8.g), 2 specifically conceived for aerial robots, and 1 for underwater robots. Most of the system are generic (47/58) and paper P40 does not provide any details about the supported robot types.

5.16. System characteristics - Platform

Another aspect characterizing robotic systems is related to the platform used for their implementation. For this parameter we consider (i) all the different frameworks that have been used in the primary study for implementation (ex. ROS, OPROS), (ii) the specific standards on top on which the platforms are based (ex. CORBA,) and (iii) tools on which they relay. Even though these platforms address different aspects and perspectives of the system and different level of abstraction (from code to architecture) we wanted to understand if there are specific frameworks used in the domain that are more common than others. While performing the analysis, we counted 17 different platforms in addition to ad-hoc ones. In Table 8 we show the most used platforms (at least two occurrences). ROS is one of the most used platforms (13/58), even though the majority of the analyzed primary studies propose their ad-hoc technologies (20/58). Such numbers are justified by the need of abstraction layers taming the complexity of writing software for robotic systems. Even though ROS was explicitly designed with such a goal, ad-hoc platforms are also employed e.g., to overcome limitations of ROS (e.g., scalability and reliability) that might be critical for some application domains.

Table 8: Platform used by the implementations of the approaches

Platform	**#Studies**	**Studies**
Ad-hoc platform	20	P6, P7, P9, P10, P13, P14, P15, P16, P19, P20, P21, P22, P25, P26, P32, P34, P42, P43, P44, P47
ROS	13	P6, P12, P17, P29, P30, P31, P32, P37, P38, P45, P51, P54, P56
OPROS	3	P29, P30, P38
ADE	2	P36, P37
Corba	2	P9, P23
OpenPRS	2	P8, P31
OrocosRTT	2	P27, P58
RTAI	2	P10, P16

5.17. Models - Model-based specification

Engineering mobile robotic systems has to take into account several aspects that might go from requirement elicitation to the specification of hardware characteristics. Consequently, the adoption of model-based techniques can help developers in managing the different aspects by increasing abstraction and enabling automation. Many approaches make use of models (42/58 as shown in Figure 9.a) for various purposes, e.g., to support the specification of missions, safety constraints, hardware invariants, etc. Only 10 approaches do not make use of models for developing and using the robotic systems at hand.

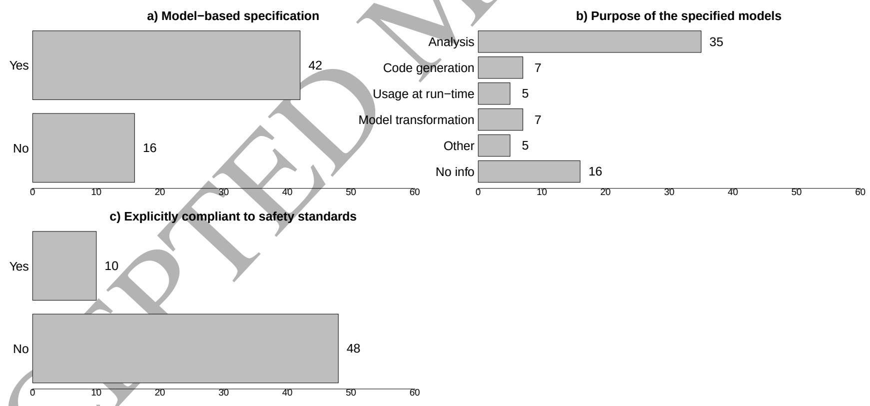

Figure 9: Model-based specifications and standards - results

5.18. Models - Purpose of the specified models

By continuing the discussion related to the previous point, the adoption of models can be done for different purposes. Most of the considered approaches (35/42 as shown in Figure 9.b) adopt models for analysis purposes (e.g., feasibility assessment, mission execution time prediction, etc.). Some of them (7/42) use models for generating the code of the modeled systems or to apply model-to-model transformations (7/42) targeting models that are in the form, which is more convenient for the particular analysis task. Some of the analyzed systems (5/42) use models at run-time e.g., to support the execution of the mission while it is executed. The papers in the *Other* category are P15, P18, P29, P37, P41. In P15 models are used to support the run-time and dynamic adaptation of systems due to unforeseen environment changes. Adaptive systems are considered also in P18 and P29 that propose the adoption of models to deal with fault tolerant aspects of the systems being developed. Fault management is also the main topic of P37, which adopts models for specifying systems consisting of multiple mobile robots. P41 proposes the adoption of models for supporting the development of autonomous systems, which have to be self-healing.

5.19. Standards - Compliant standards

Mobile robotic systems are very complex as testified also by the number of standards that are considered when developing them (see Table 9). According to Figure 9.c 10/58 approaches are compliant to standards that specifically target safety aspects. As shown in Table 9, each approach can adopt more than one standard depending on the peculiar aspects of the system being developed. For instance, P35 and P42 make use of 4 standards each. The former, proposes an approach to develop safe control systems and as a such it refers to the following standards:

- IEC61508 – Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems;
- ISO10218 – Robots and robotic devices - Safety requirements for industrial robots;
- ISO13855 – Safety of machinery - Positioning of safeguards with respect to the approach speeds of parts of the human body;
- ANSI/RIA R15.06 – Industrial Robots and Robot Systems - Safety Requirements.

In P42 authors propose an approach to verify the correctness of vision pipelines in agricultural settings with the aim of improving the safety of the systems being developed. The proposed approach considers the following standards:

- ISO13482 – Robots and robotic devices - Safety requirements for personal care robots;
- ISO25119 – Tractors and machinery for agriculture and forestry Safety-related parts of control systems;
- ISO18497 – Agricultural machinery and tractors – Safety of highly automated agricultural machines;

Table 9: Standards with compliant approaches

Standard	#Studies	Studies	Domain	Focus
IEC61508	4	P5, P21, P35, P52	Generic	Functional safety
ISO13482	3	P33, P39, P42	Personal care robots	Safety requirements
ISO10218	1	P35	Industrial robots	Safety requirements
ISO13855	1	P35	Industrial robots	Positioning of safeguards
ANSI/RIA R15.06	1	P35	Industrial robots	Safety requirements
ISO14121	1	P21	Generic	Risk assessment
ISO11199	1	P21	Generic	Requirements and test methods
ISO12100	1	P39	Generic	Risk assessment and risk reduction
IEC60204	1	P39	Generic	Electrical equipment of machines
ISO25119	1	P42	Agriculture	Safety-related parts of control systems
ISO18497	1	P42	Agriculture	Design safety principles
IEC61496	1	P42	Generic	Safety of electro-sensitive protective equipment
ISO62262	1	P9	Generic	Protection provided by enclosures for electrical equipment against external mechanical impacts
IEC61608	1	P9	Generic	Functional safety
OASIS	1	P49	Generic	Information society
RTCADO178C	1	P45	Aviation	Airborne Systems and Equipment Certification

- IEC61496 – Safety of machinery - Electro-sensitive protective equipment.

As it is possible to notice, the standards that are referred by the existing approaches vary much depend on the particular application domains where the considered robotic systems will be employed.

5.20. Hazards - Unexpected environment hazards

In order to employ mobile robotic systems in real contexts, it is important that they have the capability of reacting to unexpected environment threats, such as the presence of unpredicted obstacles, the presence of humans in the operating area, etc. We define hazard as an atomic event, situation, and/or object that brings an unavoidable danger or risk in mobile robotic systems.

Hazards can have a variety of forms (ex. an internal fault of a robot, an unwanted human behavior, an unexpected situation - dynamic obstacle, an emergent behaviour raised from the cooperation and the coordination of the robots and much more other situations coming internally from the system or externally from the environment). As shown in Figure 10.a, the majority of the analyzed systems (29/58) implement such a capability. The primary studies P3, P4, P8, and P38 do not give explicit information about that. In particular, P3 proposes an approach to support the diagnosis of complex systems. P4 discusses all the concepts that have to be taken into account when designing autonomous systems by touching different peculiar aspects like communication, control, and navigation. The focus of P8 is supporting testing activities when developing the control software for autonomous systems. With the aim of improving the quality of the software of robotic systems, P38 proposes an approach to manage faults of components based on the OPRoS platform.

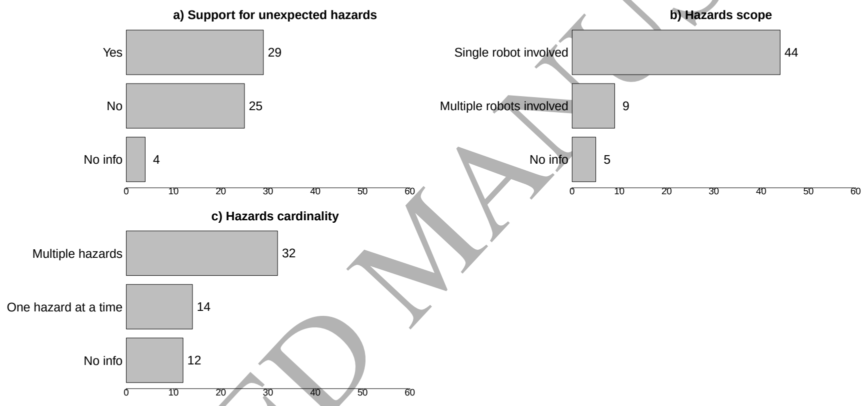

Figure 10: Hazards - results

5.21. Hazards - Scope

When considering unexpected environment hazards, systems can be distinguished with respect to their capability of managing threats impacting or due to a single robot (44/58 as according to Figure 10.b), from those occurring because of the cooperation and coordination of different robots. Only 9 out of 58 analyzed systems are able to manage unexpected hazards coming from multi-robots systems.

5.22. Hazards - Cardinality

Another level of complexity related to the management of unexpected environment hazards is related to the capability of the system to manage one or multiple threats at a time. According to the performed analyzed and as shown in Figure 10.c, most of the analyzed systems are able to deal with multiple hazards, whereas 14 out of 58 have the capability of managing only one hazard at a time. Unfortunately, 12 primary studies do not provide explicit information about such characteristic.

Highlights - Management of safety of mobile robotic systems (RQ1):

A first contribution that we obtained when answering this research question is a classification framework for identifying how safety is managed. The classification framework is graphically represented as a three-like structure in Figure 5. This figure highlights the aspects that, according to our primary studies, developers should consider when engineering a safety management solution. Here in the following we summarize these aspects.

The majority of the primary studies propose new (mainly generic) methods for achieving safety for MRSs.

The vast majority of the primary studies manage safety by relying on knowledge which is: (i) local to each robot and (ii) exploited to implement local safety mechanisms without any cooperation with other robots. Some insights about the different methods for achieving safety might be found in Figure 11.

Safety is considered at different levels of abstraction, by spanning from requirement specification till implementation, even though most of the approaches work at design level by making use of different kinds of models. Safety-specific concerns are typically specified in a platform- and robot-independent manner. Contrariwise, the actual management of safety is not kept separated from the functional management of robots.

Most of the primary studies do not seem to address safety in case of different kinds of robots and of dynamic additions or removals of robots and/or other agents. Context awareness is instead implemented by the vast majority of the analysed studies, which are able to sense some key properties of the considered operational context of robots, and consequently to implement adaptiveness capabilities in case of context changes.

Few primary studies are able to manage safety for multi-robot systems and the majority of the analysed approaches work atop of ad-hoc platforms, even though ROS is gaining more and more momentum.

Further research is still needed to overcome important limitations of MRSs, in particular the capability of reacting to unexpected environment hazards by still keeping safety under control.

Developers of safety solutions might use the framework to select the technique or the approach that better matches the characteristics of their system, as well as the nature of their hazards, etc.

How safety is managed across application fields. Mobile robotic systems is a wide domain with many specific fields, such as exploration missions, service robotics, self-driving vehicles. Totally different approaches can be applied for solving concerns that are specific for each application field. In order to provide guidance to researchers and practitioners on which application fields have been concretely investigated by researchers, in Table 10 we report the application fields which have been considered during the evaluation of the proposed approaches. Practitioners can consider this table as an indication of research approaches that can be potentially applied in real-world projects in specific application fields.

We investigated whether the application field in which a given approach has been evaluated actually correlates with specific characteristics of the approach itself (e.g., do approaches evaluated in the context of exploration missions manage self-adaptation in the same way as approaches evaluated in the medical care field?). To this goal, we analyzed the extracted data to explore the possible relation between the *application field* and all the parameters considered when answering RQ1 (e.g., *openness*, *context awareness*, *cardinality of hazards*). This results in 19 pairs of parameters, where the first one is always *application field* and the second one is one of the parameters we considered in RQ1; for each pair, we built a contingency table and evaluated the actual existence of possible relations. In the following we report the main results of our analysis.

Table 10: Recurrent application fields

Application fields	**#Studies**	**Studies**
Exploration mission	11	P6, P8, P18, P19, P22, P24, P27, P28, P34, P43, P53
Service robotics	11	P2, P10, P16, P32, P33, P34, P35, P39, P47, P50, P53
Not specified	9	P4, P11, P12, P44, P46, P48, P49, P51, P55
Search and rescue	6	P1, P5, P33, P34, P36, P41
Navigation tasks	4	P30, P37, P38, P58
Self-driving vehicles	4	P13, P14, P15, P26
Medical care	3	P7, P9, P21
Playing soccer (RoboCup)	2	P3, P23
Industrial robotics	2	P54, P57
Automatic cleaning	1	P25
Scientific research	1	P31
Transportation	1	P52
Waste cleanup	1	P2
Drawing lines of soccer fields	1	P56
Environment protection	1	P17

For what concerns the **safety management**, the majority of the approaches relies on a local knowledge of the environment, with the only exceptions of search&rescue (3 approaches), service robotics, waste cleanup (which rely on cooperative world knowledge), and industrial robots, (which rely on a centralized world knowledge).

We noticed a similar trend when considering also the scope of the safety mechanisms (i.e., local vs cooperative vs centralized), again with two exceptions (waste cleanup and service robots relying on cooperative mechanisms). When dealing with the considered abstraction levels (e.g., architecture, low-level design, etc.), we see a tendency aligned with the results of the vertical analysis (i.e., strong preponderance of low-level design), where architecture is more considered when dealing with exploration missions and navigation tasks; interestingly, requirements are more considered in the medical care and search&rescue application fields (2 approaches each). We can trace the usage of requirements in the medical domain to the need of certifications and standard compliance. The aspects related to separation of concerns, platform-independent specification, and additional property types follow the same trends as their corresponding vertical analyses.

When considering the **characteristics** of the proposed approaches, we report that openness, context awareness, and types, heterogeneity, and cardinality of robots do not exhibit strong trends with respect to the application field in which they have been evaluated. The same applies for the other parameters related to the characteristics of the approaches, but with two notable exceptions. Firstly, approaches with adaptiveness capabilities have been mostly evaluated in the context of generic robots (i.e., the evaluation has been carried out at an abstract level), robots performing navigation tasks, self-driving vehicles, and service robots. A different tendency has been observed when considering UAVs, where they have been always evaluated in the context of approaches without adaptiveness capabilities. This might be a result from the safety-criticality of the domain. UAVs are part of a domain where strong safety regulations are needed to be used in everyday life. They have a large variability space, so, in many cases, guaranteeing their safety might be a complex and intractable process. Hence, we interpret this result in the following way: most of the approaches that ensure safety for UAVs focus on safety by construction, omitting adaptiveness capabilities. This can lead to the conclusion that UAVs safety is mostly addressed at design-time. Secondly, all approaches evaluated in the context of self-driving vehicles are based on ad-hoc platforms. This can lead to the conclusion that self-driving vehicles lack a standardized platform, processes and tools for designing and analysing safety approaches. Furthermore, it is difficult to compare the different approaches across a variety of environments. We interpret the last observation as a clear indication of the need for standardization of safety-related aspects in the field of self-driving vehicles, ranging from its hardware, software, and communication perspectives.

In the context of **model-based** approaches, we observed trends aligned with the ones resulting from the vertical analysis, both in terms of being model-based and the purposes of the considered models. The only strong exception is related to the fact that service robotics have been evaluated mostly in non-model-based approaches.

No surprising trends have been discovered when dealing with **standard** compliance; we can trace this absence of trends to the low number of primary studies conforming to safety standards.

Finally, **hazards** management does not exhibit strong correlations with the application field in which the approaches have been evaluated. The only exception is related to unexpected environment hazards, which have been notably considered in the context of service robotics, medical care, and exploration robots.

Classification of the primary studies with respect to parameters of other secondary studies. To complement our classification framework, we considered other secondary studies [34, 35, 36] that are somehow related to our work but that defined the parameters for managing safety in a more top down approach, instead of extracting these parameters from the considered primary studies. We then classified the analysed primary studies with respect to the parameters identified in these secondary studies. The parameters are described in Table 3 and the results of the classification is summarized in the heatmap shown in Figure 11.

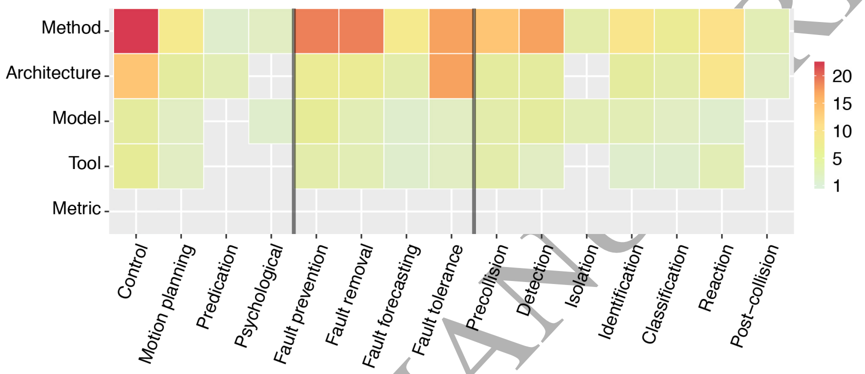

Figure 11: Classification of the primary studies with respect to parameters of other secondary studies

It is important to notice that for each of these parameters we just report a binary variable assessing whether the parameter is evaluated positively or otherwise. For what concerns the first 4 parameters, i.e. the ones coming from [36], control is the most used ones (29 approaches out of 58) followed by planning (12 out of 58), predication (4 out of 58), and finally psychological (2 out of 58). We also crossed-tabulated the results with the types of research contribution in Table 5. As it is visible in the heatmap in the figure, most of the approaches propose a method and then an architecture.

For what concerns the other four parameters, the ones coming from [34], many approaches support fault tolerance (28 out of 58), fault prevention (24 out 58), and fault removal (23 out of 58). Few approaches support fault forecasting (11 out of 58). Most of the approaches propose methods and interestingly, architecture solutions are popular for what concerns fault tolerance.

For what concerns the remaining seven parameters coming from [35], the most common solutions are into precollision (18 of 58), detection (18 out of 58), reaction (15 out 58), and identification (12 out of 58). Few are into classification (9 out of 58), isolation (4 out 58), and postcollision (4 out 58). Again no surprises here, most of the approaches propose methods and some architectures. The remaining research contributions are not very representative.

A complete description of the parameters identified in the other secondary studies and the raw data we extracted for each of them are available in the replication package of this study.

6. Potential for industrial adoption (RQ2)

In this section we will discuss the results on how existing research on safety for mobile robotic systems can be potentially adopted in real industrial projects.

6.1. Applied research method

As discussed in [7] and [49], from a high-level perspective a research solution can be assessed by means of two main research methods: *validation* and *evaluation*. Concretely, *validation* focuses on specific properties of the proposed solution and it is done in a controlled setting or in the lab; *evaluation* aims at investigating on the new situation brought by the proposed solution and it takes place in real-world (industrial) contexts. In the context of this study, evaluation potentially provides a higher level of evidence about the practical applicability of a proposed approach for safety of mobile robotic systems.

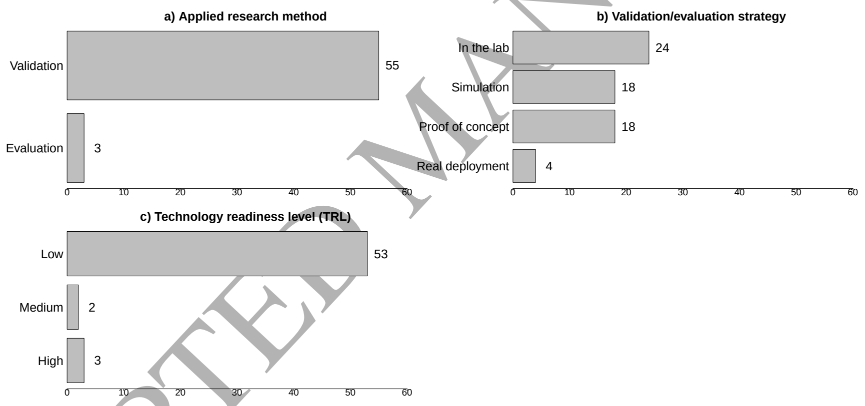

Figure 12: Applied research method (a), validation/evaluation strategy (b), and technology readiness level (c) - results

As shown in Figure 12.a, the vast majority of our primary studies provides only a *validation* of the proposed approach (55/58). This result is a clear call for researchers on safety for mobile robotic system for assessing their approaches on real-world industrial contexts, potentially leading to a smoother technology transfer of their proposed research. As a starting point for achieving this result we can get inspired by the three primary studies presenting a thorough *evaluation* of the proposed approach, they are briefly discussed in the following:

- P17 - The goal of this approach is to avoid failures of a ROS-based robotic system under various scenarios. By starting from a known training set, it automatically performs inference and monitoring of specialized invariants during the lifetime of the system. The approach has been evaluated in the context of two case studies. The first case study is about a real UAV (i.e., an Ascending Technologies Hummingbird) landing on a moving platform (realized as an iRobot Create with a mounted landing platform) under different scenarios (e.g., normal, wind blowing, fragile platform, occupied platform, false airport), whereas the second case study is about a water sampling UAV, where a combination of ultrasonic, air-pressure, GPS, and conductivity sensors are used.

- P54 - This approach makes use of model-based testing and diagnosis for supporting the dependability of autonomous robots along the whole life cycle. The approach has been evaluated in the context of a real industrial installation of autonomous transport robots in a warehouse; the system includes a fleet of individual autonomous robots, a conveyor for transportation, and a central station.
- P57 - This approach presents HAZOP-UML, a method for the safety analysis of human-robot interaction; the method supports safety analysts in specifying dynamic models of the system in UML, and in identifying hazards, recommendations, and hypotheses of possible deviations of the the system from the specified dynamic models. The approach has been evaluated by recruiting professional safety analysts and letting them apply the proposed approach on three different case studies involving (i) an assistive robot for the autonomous movement of the elderly, (ii) a KUKA Omnirob mobile robot with a KUKA Light Weight Robot arm used in workshops or factories with human workers, and (iii) a custom robot capable of navigating autonomously within a manufacturing setting while avoiding human workers, and taking and placing part boxes either on shelves or on its own base.

6.2. Validation/evaluation strategies

The analyzed studies apply different strategies for assessing their proposed approaches, independently of whether they are performed in the context of validation or evaluation research. More specifically, our analysis revealed the following assessment strategies (in order of potential realism): (i) proof of concept implementation running on simple examples, (ii) simulation-based execution and experimentation of the system, (iii) laboratory experiment where real robots are used but in a controlled environment, and (iv) realized system deployed and running in real environment.

As shown in Figure 12.b, the majority of the studies assess the approach in the lab (18/58), followed by proof of concept and simulation-based validations (18/58), and experiments on real deployments (4/58).

It goes without saying that validating research results in a real deployment is the best case in terms of potential for industrial adoption, and the authors of 4 studies managed to achieve this very ambitious goal (P17, P45, P54, P57).

Nevertheless, we have also to acknowledge that in some cases this kind of strategy is not practical if not feasible, for example in large-scale systems involving safety issues (e.g., a fleet of flying drones in a tactical environment). These are the situations were laboratory experiments may be performed in a more manageable manner. Also, it is important to say that recently simulation environments are gaining a lot of attention thanks to the great advances they are making in terms of realism of the simulation, configurability, and possibility to run software- or hardware-in the loop simulations. The latter are enabled by the high level of decoupling provided by platforms or communication middleware like ROS, where engineers can use the same software stack as the one used in real deployments, while simulating only the components depending on the real world (e.g., drivers for the GPS, accelerometers).

The relatively high number of strategies based on proofs of concept (18/58) is somehow disappointing, specially in light of todays wide availability of software platforms, simulators, and low-cost hardware components. Assessing a scientific result via a simple proof of concept and an example is not acceptable anymore in our research community. We expect that in future researchers on safety for mobile robots will move on from this comfort zone and will start providing more tangible (empirical) results and benchmarks about the performance of their proposed solutions. This will surely boost the potential for industrial adoption of our research.

6.3. Technology readiness level (TRL)

The purpose of the TRL is to objectively assess the maturity of a particular technology [50] on a scale ranging from 1 (minimum) to 9 (maximum). In order to keep the data extraction activity manageable and less time consuming, in the context of this work we classify the TRL of each primary study on a 3-levels scale: (i) *low* TRL (i.e., TRL ≤ 4), where a technology is either formulated, validated or demonstrated at most in lab, (ii) *medium* TRL (i.e., $5 \leq$ TRL ≤ 6), where a technology is either validated or demonstrated in a relevant environment, and (iii) *high* TRL (i.e., TRL ≥ 7), where the technology is either completed, demonstrated, or proven in operational environment.

Figure 12.c shows the distribution of the TRL levels of our primary studies. The obtained results are self explicative, the majority of approaches (53/58) have a low readiness level, whereas only two of them are in the medium (P51, P53) and high (P17, P54, P57) levels of TRL. This is a confirmation of the results about the evaluation and validation strategies; again, if we aim at making our research products adoptable by industry, we will need to work on their technological readiness with well tested and designed tools, and realistic experimentation.

6.4. Rigor and Industrial Relevance

As discussed in Section 3.3, we extracted data related to rigor and industrial relevance of the primary studies by applying the well-defined classification model introduced by Ivarsson et al. [37]. Specifically, we (i) read in details each primary

study, with a special focus on the sections related to the evaluation of the proposed approach, (ii) assigned a score to each criteria related to rigor and industrial relevance by carefully applying the scoring rubric proposed in [37, §3], and (iii) identified outliers in terms of total scores and manually checking and discussing them in order to identify possible errors in the score assignments. This activity has been performed iteratively by two researchers in collaboration, with the help of a third one in case of conflicts or unclear situations.
Rigor is defined as the precision, exactness, or correctness of use of the research method applied in a scientific work [37]. Intuitively, an experiment reported in such a way that its operational context is defined, its design is clear, and its threats to validity are explicitly discussed has higher rigor than an informal description of a running example. The main rationale for considering rigor in our research is that a primary study with high rigor is easier and more straightforward to be assessed by practitioners. Based on [37], the rigor of each primary study has been assessed according to the criteria in Table 11, where each criteria can be scored with the following score levels: strong (1 point), medium (0.5 point), weak (0 points). Thus, a primary study can have a rigor score ranging from 0 to 3.

Table 11: Rigor assessment criteria [37]

Criteria	**Description**
Context	Is the context described to the degree where a reader can understand and compare it to another context?
Study design	Is the study design described to the degree where a reader can understand its main parts, e.g., variables, treatments, etc.?
Validity	Is the validity and threats of the study discussed and measured in details?

The upper part of Figure 13 shows how the considered primary studies are distributed in terms of total rigor score. Here we can notice that the majority of primary studies (42/58) have a score between 0.5 and 1.5, with a mean of 1.27. Also, only 5 studies have a rigor score above 2 (P2, P31, P33, P44, P47). This result is already quite interesting: it clearly shows that researchers on safety for mobile robots should improve in terms of rigor (e.g., precision and correctness) when evaluating their research results. It also means that the majority of evaluations performed in our primary studies are either (i) experiments where rigor-related aspects are poorly reported or (ii) simple applications of the proposed approaches to toy examples. This is a clear call to researchers in the field to both better report their experiments and to focus on key aspects of the proposed approach (e.g., managed hazards, types and quality of safety-related solutions), rather than simply illustrating its application to an example.

In order to better understand this phenomenon, we dig into the scores of all the criteria for rigor of evaluation. As shown in the lower part of Figure 13, the context and the study design are performing quite well, with the majority of studies falling within the medium/strong score levels. The real problem with

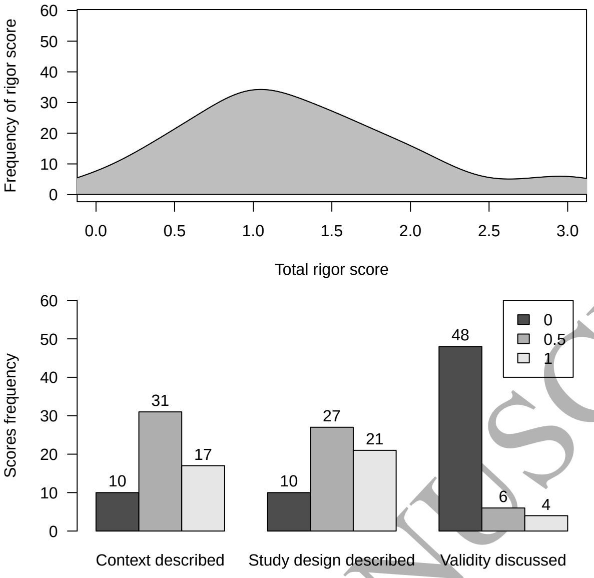

Figure 13: Results for rigor of evaluation

rigor lies in the identification and reporting of the validity of the performed evaluations; indeed during this research we seldom noticed that the threats to validity of the performed experiment have been thoroughly discussed. To understand if the data extracted from the primary studies really reflects the conclusion and results of the authors, we contacted the first authors of each primary study and we incorporated their comments in our findings. Of course, understanding how valid the results of an evaluation/experiment are is a fundamental aspect for the adoptability of a proposed approach. As a solution, we suggest researchers to carefully consider all the potential threats to validity of their performed evaluations and to explicitly report them; as suggested in [9], this activity should be already carried out in the planning phase of an evaluation/experiment. Also, for easing the design, understanding and replicability of the performed evaluations, it is suggested to structure the discussion of threats to validity according to well-known classification schemes, such as the one by Cook and Campbell [51].

Industrial relevance refers to the realism of the evaluation of an approach, and determines the potential relevance of its results for industry [37]. Intuitively, an experiment involving a large number of professionals as subjects and deploying the robots in a real operational environment has a higher industrial relevance with respect to a software simulation performed in a research lab.

Table 12 shows the criteria we used for assessing the industrial relevance of each primary study. In conformance with [37], we assessed a primary study for each industrial relevance criterion as either strong (1 point) or weak (0 points). A primary study can have an industrial relevance score ranging from 0 to 4.

Table 12: Industrial relevance assessment criteria [37]

Criteria	**Description**
Subjects	Are subjects used in the evaluation representative (real robots)?
Context	Is the evaluation performed in a representative setting (e.g., real deployment environment)?
Scale	Is the scale of the applications used in the evaluation of realistic size (e.g., size of the operational environment, number of involved robots)?
Research method	Does the applied research method facilitate the investigation of real situations (e.g., an industrial case study)?

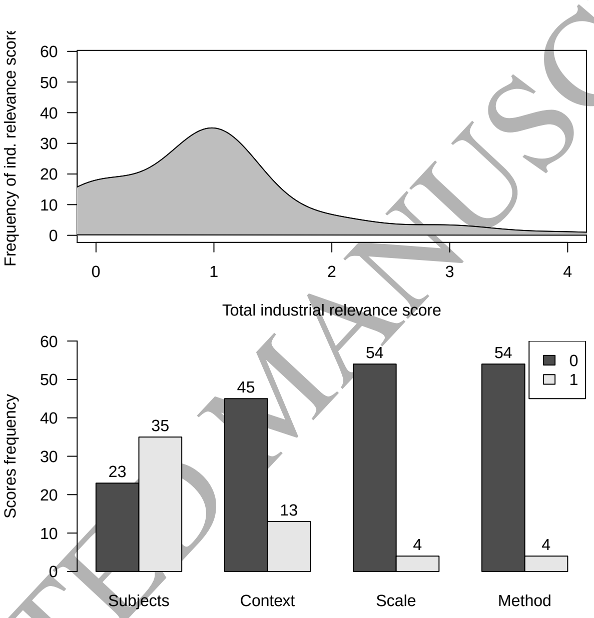

Figure 14: Results for industrial relevance

Similarly to the rigor score, the distribution of the primary studies with respect to their total industrial relevance score is not showing good results. Indeed, referring to the upper part of Figure 14, the majority of primary studies (54/58) scores lower than 2. If we zoom into the specific criteria, in the lower part of Figure 14 we can notice that research on safety for mobile robots suffers in terms of the context, scale, and research method dimensions. More specifically, it emerged that almost all primary studies do not report on the evaluation of the proposed approach in a representative setting (*context* criterion, 45 studies), with a realistic size (*scale* criterion, 54 studies), or facilitating a real investigation (*research* method, 54 studies). It is important to point out that we are not evaluating the validity of an approach in this way, but these are all aspects that researchers should take into consideration if their aim is to develop methods that should be adopted in real industrial settings. On a positive side, the *subjects* score has a good performance. Researchers achieved this result by using in many cases real robots for their evaluations. This can be seen as a consequence of opportunities opened by open software/hardware platforms for robotics, making them accessible at low prices and with the needed level of configurability.

6.5. Industry involvement

In this section we aim at characterizing the involvement of practitioners into research studies on safety for mobile robots. Inspired by the classification used in a previous work [52], we categorize each primary study as: *academic* if all authors are affiliated with universities or research centers, *industrial* if all authors are affiliated with some companies, or a *mix* of the previous two categories. It is important to note that the number of involved industrial authors can be considered as an upper bound, as an industrial affiliation does not strictly mean that industry was actively involved in the performed research.

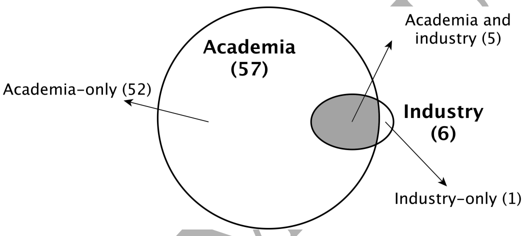

Figure 15: Distribution of industry involvement

Figure 15 shows that almost all primary studies contribute with an academic-only perspective (52/58). Then, only 5 studies contribute with a mixed perspective and only one study provides an industry-only perspective. This result is somehow aligned with the analysis of the previous aspects and it clearly shows the low involvement of industrial partners in research on safety for mobile robotic systems. This result is a sign of a missed opportunity; research on this research area seems to have been performed in isolation with respect to the industrial perspective, which may bring new relevant problems to be solved and a much clear picture of the state of the practice in the field. Researchers and practitioners on safety for mobile robots should work together on creating better synergies and cooperation plans so that research will be performed on industrially relevant problems and new research methods, technologies and tools will smoothly transition from academia to industry [9].

Highlights - Industrial adoption of existing approaches for safety of MRSs (RQ2)

The technology readiness level showed that most of the approaches are not mature enough to be used in real industry settings. Most of the primary studies validated the proposed approaches in the lab and very few considered real deployments. This can be enough as a proof of concept, however, more work is needed in order to use these approaches in robotic applications that are supposed to be used in real environments. Moreover, most of the approaches do not provide an identification and reporting of the validity of the performed evaluations; indeed during this research we seldom noticed that the threats to validity of the performed experiment have been thoroughly discussed. In other words, important efforts have to be spent to transfer the approaches, that currently were validated by means of proof-of-concept implementations, to real-world industrial contexts. According to our experience, an effective way to reach this objective is to have a more significant involvement of industrial partners in the development and validation of techniques and approaches for the management of safety for MRSs. Additionally, the involvement of industrial partners is only a necessary but not sufficient condition for successfully transferring a research product into industry; at least setting up a proper documentation, tool support, and a concrete knowledge transfer plan are evenly important activities, which should be proactively pursued by researchers.

7. Emerging challenges on safety for MSRs (RQ3)

In this section we discuss the main findings of the paper as well as their implications for future research.

7.1. Single vs Multi-robots

As discussed in Section 5.14, most of the primary studies focus on a single robot (45/58). We acknowledge that there is the need of solutions to manage safety at the level of single robot, however, there is also the need of approaches that deal with multiple robots. In fact, the collaborative smart robots market size is expected to reach USD 1.07 billion by 2020 whereas the software market size for smart robots is expected to grow at a CAGR of 30.24% from 2015 to 2020 [53].

As implication for future research, we highlight the need of solutions addressing safety when multiple robots need to collaborate with each other in order to accomplish complex missions. These approaches might require cooperative safety management mechanisms (see Section 5.4) and cooperative or centralized world knowledge (see Section 5.3).

7.2. Openness and capability to cope with uncertainty

In the near future, MRSs will be used in tasks of everyday life. This means that often MRSs will be used in unknown or partially unknown environments that might be shared with humans or other robots. This will require context awareness, and most of the approaches in our primary studies (41/58) have these capabilities (see Section 5.10), and adaptiveness capabilities to changing environments. As shown in Section 5.11, 25/58 approaches do not support adaptiveness capabilities and 4 approaches provide no information. Moreover, as shown in Section 5.9, only 5 approaches out of 58 are able to deal with open systems, meaning that in those cases new robots or human actors can be added at runtime.

As implication for future research, the adoption of MRSs in tasks of everyday life will require more investigation in adaptiveness capabilities as well as in dealing with open systems. In the case of MRSs will need to deal with partially known and uncontrollable environments, machine learning seems to be a promising solution that is getting increasing attention. However, the use of machine learning in safety-critical domains is still an open problem and innovative solutions are needed. A promising approach is to combine machine learning with run-time verification techniques [48].

7.3. Compliance to standards

MRSs are very complex systems and consequently advanced techniques and tools are needed for supporting their development. Especially for critical systems, safety represents a crucial aspect to be managed since the early stages of development. In this respect, over the last decade several standards have been issued to manage MRSs safety. As shown in Table 9, dozens of standards are available for safety. Each application domain has its own specificities and this might justify the need of dedicated standards. Following this reasoning, in the future we might have the definition of further standards due to the increasing adoption of robotic systems in different application scenarios.

7.4. Adoption of model-driven engineering for robotic systems

Model-Driven Engineering refers to the systematic use of models as first-class entities throughout the software engineering life cycle. The objective is to increase productivity and reduce time to market by enabling the development of complex systems by means of models defined with concepts that are much less bound to the underlying implementation technology and are much closer to the problem domain. According to our study, models are currently used in the domain of robotic systems for different purposes e.g., to support the specification of missions to be executed by robots, safety constraints, etc. Most of the analysed approaches (38 out of 48) take advantage of models for performing analysis tasks since the early stages of development. This is justified by the fact that design-time models help the understanding of complex problems and their potential solutions through abstractions. As also envisioned by the Robotics 2020 – Multi-Annual Roadmap[8], for the future we foresee the exploitation of run-time models, which will be used to support monitoring and diagnosis of robots, to explain what robots are doing during the execution of defined missions, and even to perform dynamic adaptations that might occur after MSR missions are started. To this end, the main challenge that should be investigated in the future is the proper management of MSR run-time models and the possibility to trace them back to design ones. For instance, the MegaM@aRt2 EU

---

[8] https://www.eu-robotics.net/cms/upload/downloads/ppp-documents/Multi-Annual_Roadmap2020_ICT-24_Rev_B_full.pdf

ECSEL project[9] is conceiving techniques and tools for supporting the traceability across different layers of complex cyber-physical systems ranging from highly specialized engineering design models to low-level log entries. Traceability tools are being developed in the project in order to preserve and exploit traceability information between different layers of abstraction, notably to provide developers with reusable feedback from runtime to design time. Thus a methodological loop (supported by megamodeling and model transformation techniques) between models at design-time and run-time levels is under investigation in the MegaM@aRt2 project with the final aim of supporting model-based continuous development and validation of large and complex systems [54].

7.5. Rigor and Industrial Relevance

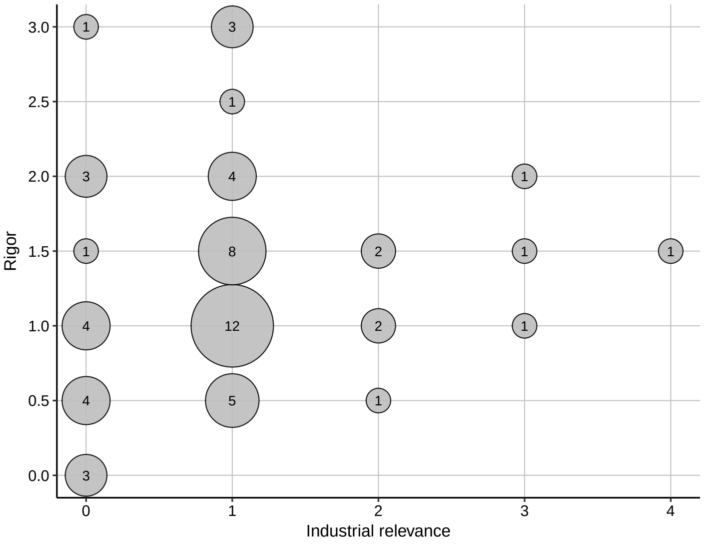

Figure 16: Aggregation of scores for rigor and industrial relevance

As discussed in Section 6.4, the majority of evaluations in safety for robotic systems lack both rigor and relevance. This result is even more evident when considering these two dimensions together. The bubble chart in Figure 16 graphically shows the aggregation of rigor and industrial relevance of the primary studies. Here the majority of the primary studies falls in the lower-left quadrant, highlighting the lack of both rigor and industrial relevance. Given the situation, in the following we propose a set of strategies for improving the evaluation of robotic systems in terms of both rigor (i.e., moving ↑ in Figure 16) and industrial relevance (i.e., moving → in Figure 16):

---

[9] https://megamart2-ecsel.eu/

- improve the design of the performed experiments by, e.g., formalizing the safety hazards being considered, explicitly defining the dependent/independent variables of their experiments, identifying sound statistical analyses of the obtained data (↑);
- elaborate on and discuss potential threats to validity before and after evaluating the robotic system (↑);
- improve the measurement precision when performing experiments involving both the software and hardware parts of the robots (↑);
- carefully select the software and hardware platforms for the evaluation, preferably using real robots (→);
- carefully select realistic operational environments where the robots will be deployed (→);
- push towards large-scale, or at least realistic-scale evaluations, involving a realistic number of robots and involved human users, this is especially true for swarm and multi-robot systems (→);
- when possible, push towards investigating real situations involving industrial partners, practitioners, and in-the-field operators (→).

Researchers can use the above mentioned strategies to ensure an adequate rigor and relevance when planning the evaluations of approaches for safety of robotic systems.

7.6. Software engineering and robotics

As stated by the H2020 Multi-Annual Robotics Roadmap ICT-2016 [55], in the production of software for robotic systems "usually there are no system development processes (highlighted by a lack of overall architectural models and methods). This results in the need for craftsmanship in building robotic systems instead of following established engineering processes." The use of ad-hoc development processes in general, and software engineering approaches in particular, hampers reuse and complicates the configurability of existing solutions. This justifies the need of systematic approaches, methods, and tools to (i) easily configure robots, or provide them with self-configuration capabilities, (ii) specify robotic tasks in an easy and user-friendly way, and (iii) make the robots able to take decisions on their own to manage unpredictable situations. This shifts towards well-defined engineering approaches will stimulate component supply-chains and significantly impact the robotics marketplace.

Even though there is a growing interest (see Section 4), the community of software engineering and robotic is still not consolidated. This is testified by the extreme fragmentation of the targeted publication venues, as discussed in Section 4. There are some workshops and initiatives in the direction of creating a community around software engineering and robotics, such as the International Workshop on Robotics Software Engineering (RoSE'18), colocated with ICSE2018, which attracted at the first edition more than 30 participants, the international workshop on Domain-Specific Languages and Models for Robotic Systems (DSLRob), the series of workshops on Model-Driven Robot Software Engineering (MORSE), the Journal of Software Engineering for Robotics (Joser), the International Conference on Robotic Computing (IRC) and a recent technical briefing at ICSE on software engineering for robotic systems [56]. However, more work is needed in order to create a proper community on this topic.

Highlights - What are the main emerging challenges for future research on safety for mobile robotics systems? (RQ3)

We found that most of the approaches surveyed in this study focus on a single robot. Therefore, when multiple robots need to collaborate each other in order to accomplish complex missions, it emerges then the need of solutions addressing safety for MRSs.

Many of the surveyed approaches do not support adaptiveness capabilities and most of them are not able to deal with systems supporting the addition and removal of robots, human actors, etc. at runtime. Tasks of everyday life will require more investigation in safety-oriented adaptiveness capabilities of MRSs.

Many domain-specific standards related to safety are currently available. However, only a minority of the surveyed approaches are compliant to standards targeting safety aspects. Consequently, when developing a robotic system, specific standards have to be taken into account to make it compliant to them and safe for the considered application domain.

The majority of evaluations in safety for robotic systems lack both rigor and relevance. Therefore, there is the need of new strategies to better support and planning the evaluations of approaches for safety of robotic systems.

Even though there is a growing interest and some relevant initiatives, the community of Software Engineering for robotics is still not consolidated. The challenge for the research community is to promote a shift towards well-defined engineering approaches able to stimulate component supply-chains and significantly impact the robotics marketplace.

8. Threats to Validity

The quality of our research has been ensured by defining a complete research protocol beforehand, by letting it assess by independent reviewers, and by conducting research following well-accepted guidelines of systematic review/mapping study [8, 7, 9]. Also, to allow independent replication and verification of our study, a complete replication package is publicly available[10] to interested researchers. Our replication package includes the review protocol, the list of all

---

[10] http://cs.gssi.it/safetyMRSReplicationPackage

considered and selected studies, the description of the parameters for the data extraction activity (i.e., the data extraction form), the raw extracted data, and the R scripts for data analysis.

In the following we discuss how we considered and mitigated the potential threats to validity of our study by following the Cook and Campbell classification framework for threats to validity [9].

Conclusion validity. Conclusion validity refers to the relationship between the extracted and synthesized data and the produced map and findings [9].

In order to mitigate possible conclusion biases, first of all we sistematically defined the search string of our automatic search (see Section 3.2) and we documented all the steps of our research in a publicly available research protocol. This allows third-party researchers to replicate our study independently.

Moreover, we documented and we used a rigorously defined data extraction form, so that we have been able to reduce possible biases that may happen during the data extraction process; also, in so doing the data extraction process can be considered as consistent and relevant to our research questions.

On the same line, the classification framework may be another source of threats to the conclusion validity of our study; indeed, other researchers may identify classification frameworks with different facets and attributes. In this context, we are mitigating this bias by (i) performing an external evaluation by independent researchers who are not directly involved in our research (see Section 3, and (ii) having the data extraction process conducted by the principle researcher and validated by the secondary researcher.

Internal validity. Internal validity is concerned with the degree of control of our study design with respect to potential extraneous variables influencing the study itself.

In this case, having a rigorously defined protocol with a rigorous data extraction form helped in mitigating biases related to the internal validity of our research. Also, for what concerns the data analysis validity, the threats are minimal since we employed only descriptive statistics when dealing with quantitative data. When considering qualitative data, we systematically applied the keywording method for transforming qualitative data into quantitative data. Finally, 10 primary studies have been randomly selected and two researchers checked whether the results were consistent, independently from the researcher performing the extraction; moreover, each disagreement has been discussed and resolved, together with a third researcher, when needed.

Construct validity. Construct validity concerns the validity of extracted and synthesized data with respect to our research questions. Construct validity concerns the selection of the primary studies with respect to how they really represent the population in light of what is investigated according to the research questions.

Firstly, we are reasonably confident about the construction of the search string used in our automatic search since the used terms (e.g., safety, mobile robotic system, etc.) have been piloted in preliminary searches (using the IEEE Xplore library); also, the chosen terms of the search string have been evaluated by the reviewers of our research protocol beforehand. As described in

Section 3.2, the automatic search has been performed on multiple electronic databases to get relevant studies independently of publishers' policies and business concerns. The used electronic databases cover the area of software engineering well [25, 26], and we are reasonably confident that this applies also to safety for mobile robotic systems from the software engineering point of view. As highlighted along the entire paper, the focus of this work is on software aspects, this is why the selection of these databases is appropriate. Moreover, domains different from robotics might be relevant to study safety aspects, however, we leave these aspects out of this study since opening to other domains would bring ealsily to an intractable number of papers to be considered.

Moreover, we complemented the automatic search with the snowballing activity performed in stage 3 of our study search and selection process (see Figure 2), thus making us even more confident about the search strategy of this study. Since our automated search strategy actually relies on the quality of the used search engines and on how researchers write their abstracts, the set of primary selected studies has been extended by means of the multi-step snowballing procedure (see stage 2 in Figure 2).

After having collected all relevant studies from the automatic search, we rigorously screened them according to well-documented inclusion and exclusion criteria (see Section 3.2); this selection stage has been performed by the principle researcher, under the supervision of the secondary researcher. Also, in order to assess the quality of the selection process, both principle and secondary researchers assessed a random sample of studies, and the inter-researcher agreement has be statistically measured with good results (see Section 3.2). Because of all the above mentioned strategies for mitigating possible threats to the construct validity of our research, we are reasonably confident that we unlikely missed potentially relevant studies.

Finally, we are aware that when analyzing the potential for industrial adoption (RQ2) we focus only on the information reported in the primary studies (for example, we do not consider knowledge transfer activities/events/initiatives around each proposed research). Even though the applied research methods, TRL level, rigor, industrial relevance, and industry involvement may be good indicators for the potential for industrial adoption of a research product in robotics, in this study we are not considering other evenly important factors such as: setting up a proper documentation, pursuing a stable tool support, building a wide and motivated community, or designing an effective knowledge transfer plan. Those aspects fall outside the scope of this study and can be targeted by future studies.

External validity. It concerns the generizability of the produced map and of the discovered findings [9]. To mitigate the threat of possible misunderstanding the conclusions from the primary studies, we contacted the first authors of each primary study and presented to them our mapping study. This way we were able to confirm that the data we extracted from the primary studies reflects the authors' findings. All their comments that were in line with the direction of our paper were thoroughly discussed and incorporated.

In our research, the most severe threat related to external validity consists in having a set of primary studies that is not representative of the whole research on safety for mobile robotic systems. In order to mitigate this possible threat, we employed a search strategy consisting of both automatic search and double-step snowballing of the primary selected studies. Also, having a set of well-defined inclusion and exclusion criteria contributed to the external validity of our study.

Moreover, only studies published in the English language have be selected in our search process. This decision may result in a possible threat to validity because potentially important primary studies published in other languages have not been selected in our research. However, the English language is the most widely used language for scientific papers, so this bias can be reasonably considered as minimal.

Similarly, grey literature (e.g., white papers, not-peer-reviewed scientific publications) is not included in our research; this potential bias is intrinsic to our study design, since we want to focus exclusively on the state of the art presented in high-quality scientific papers, and thus undergoing a rigorous peer-reviewed publication process is an accepted requirement for this kind of scientific works.

9. Related Work

In this section we discuss those secondary studies which completely or partially are addressing the topic of safety in MRSs.

The authors of [20] present a general survey of various publications that focus on mechanical design and actuation, controller design and safety criteria and metrics used to validate safety of a domestic robot during unexpected collisions between a robot and a human user, without elaborating the separate papers in details. Furthermore, the focus on the survey is on the mechanical and controller design, while not taking in consideration safety from a software engineering point of view.

A review about Human-Robot Interaction (HRI) is presented in [21]. It attempts to identify the key themes and challenges from multiple perspectives, as HRI requires understanding and comprehension of multiple domains related to people, robotics, design, cognitive psychology etc.

A survey investigating safety issues in human-robot interactions is proposed in [22]. It starts with a review of safety issues in industrial settings, then shifting focus on safety issues related to mobile robots that operate in dynamic and unpredictable environments. It gives general ideas and directions of possible hazards and methods used for risk reduction, pointing out risks being introduced with the development of modern robotic systems.

[36] presents a survey of methods for safe human-robot interaction. It discusses a variety of methods ranging from physical contact to adverse psychological effects resulting from unpleasant or dangerous interaction. The works are classified into four major categories: safety through control, motion planning, prediction, and consideration of psychological factors.

The authors of [34] present survey on dependability techniques used for increasing safety in robots. The survey reviews the main issues, research work and challenges in the field of safety-critical robots, linking up concepts of dependability and robotics.

Finally, the authors of [23] present the state of the art and enlighten a number of challenges in the field of safe and dependable physical human-robot interaction undertaken within two projects: PHRIDOM (Physical Human-Robot Interaction in Anthropic Domains) and PHRIENDS (Physical Human-Robot Interaction: dependability and safety). Results from different research groups about possible metrics for the evaluation of safety, dependability and performance in physical human-robot interaction are presented. The sources for the discussion on physical human-robot interaction is based on number of articles taken from predetermined workshops, European projects and journals.

All aforementioned studies are surveys that include couple of the most important papers in a specific sub-field of the domain. This means that the works included are not representative for the overall domain considered in this study. On another note, they do not provide a systematic way for classification of the different works.

10. Conclusions

In the near future, MRSs will need to be able to operate in uncontrollable and unknown environments. Moreover, often MRSs will be required to collaborate both with each other and with humans, to accomplish complex missions. In the last decades, robotic research has made huge progresses. However, as this study testifies, existing solutions are not yet ready to be used in everyday life, and in uncontrollable and unknown environments often shared with humans. We came to this conclusion through a mapping study devoted at investigating how existing solutions for MRSs address safety aspects. Specifically, the three research questions we investigated are:

- **RQ1**: *How do existing approaches address safety for MRSs?*
- **RQ2**: *What is the potential for industrial adoption of existing approaches for safety for MRSs?*
- **RQ3**: *What are the main emerging challenges for future research on safety for mobile robotics systems?*

The classification resulting from our investigation on RQ1 provides a solid foundation for *researchers* willing to further contribute this research area with new approaches for safety MRSs, or willing to better understand or refine existing ones. Our results with respect to RQ2 can be of special interest for *practitioners* since they provide an evidence-based instrument for identifying which approaches for safety for MRSs are the most ready to be transferred to industry. By answering RQ3 we present the main challenges and implications for future research on safety for MRSs.

In summary, this study provides a comprehensive and replicable picture of the state of the art on safety for MRSs, helping researchers and practitioners in finding characteristics, limitations, and gaps of current research on safety for MRSs. We believe and we hope that the results of this study will lead to the development of new methods and techniques for safety for MRSs, making them one step closer to supporting us in our everyday tasks of the near future.

Acknowledgments

Research partly supported from the EU H2020 Research and Innovation Programme under GA No. 731869 (Co4Robots) and from the European Research Council (ERC) under GA No. 681872 (DEMIURGE).

References

[1] M. Ford, Rise of the robots: Technology and the threat of a jobless future, Organization Management Journal 13 (2) (2016) 115–117. doi:10.1080/15416518.2016.1180076.

[2] H. Kagermann, W. Wahlster, J. Helbig, Recommendations for implementing the strategic initiative industrie 4.0 – securing the future of german manufacturing industry, Final report of the industrie 4.0 working group, München (2013).
URL http://forschungsunion.de/pdf/industrie_4_0_final_report.pdf

[3] Drone investment trends 2018, Tech. rep. (2018).
URL https://www.droneii.com/drone-investment-trends-2018

[4] Unmanned aerial vehicle (uav) - forecast & analysis to 2014 - 2020, Tech. rep. (2016).
URL http://www.marketsandmarkets.com/Market-Reports/unmanned-aerial-vehicles-uav-market-662.html

[5] Unmanned aerial vehicle market by application, class, system, uav type, mode of operation, range, point of sale, mtow and region - global forecast to 2025, Tech. rep. (2018).
URL https://www.marketsandmarkets.com/Market-Reports/unmanned-aerial-vehicles-uav-market-662.html

[6] E. Mitka, A. Gasteratos, N. Kyriakoulis, S. G. Mouroutsos, Safety certification requirements for domestic robots, Safety science 50 (9) (2012) 1888–1897.

[7] K. Petersen, S. Vakkalanka, L. Kuzniarz, Guidelines for conducting systematic mapping studies in software engineering: An update, Information and Software Technology 64 (2015) 1–18.

[8] B. A. Kitchenham, S. Charters, Guidelines for performing systematic literature reviews in software engineering, Tech. Rep. EBSE-2007-01, Keele University and University of Durham (2007).

[9] C. Wohlin, P. Runeson, M. Höst, M. Ohlsson, B. Regnell, A. Wesslén, Experimentation in Software Engineering, Computer Science, Springer, 2012.

[10] B. Siciliano, O. Khatib, Springer handbook of robotics, Springer, 2008.

[11] Introduction to robots, http://www.galileo.org/robotics/intro.html (October 2014).

[12] Oxford dictionary: Origin of the term robot, https://www.oxforddictionaries.com/definition/english/robot (October 2014).

[13] T. Harris, How robots work - howstuffworks (2014).
URL http://science.howstuffworks.com/robot.htm

[14] E. Garcia, M. A. Jimenez, P. G. De Santos, M. Armada, The evolution of robotics research, Robotics & Automation Magazine, IEEE 14 (1) (2007) 90–103.

[15] Safety standards: International organization for standardization (ISO) for robots and robot systems integration, https://www.iso.org/standard/41571.html (October 2014).

[16] A. Avižienis, J.-C. Laprie, B. Randell, Dependability and its threats: a taxonomy, in: Building the Information Society, Springer, 2004, pp. 91–120.

[17] T. Skrzypietz, Unmanned Aircraft Systems for Civilian Missions, BIGS policy paper: Brandenburgisches Institut für Gesellschaft und Sicherheit, BIGS, 2012.
URL http://goo.gl/KD8x2k

[18] O. Ogorodnikova, How safe the human-robot coexistence is? theoretical presentation, Acta Polytechnica Hungarica 6 (4) (2009) 51–74.

[19] Y. Nakabo, H. Saito, T. Ogure, S. H. Jeong, Y. Yamada, Development of a safety module for robots sharing workspace with humans, in: Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on, IEEE, 2009, pp. 5345–5349.

[20] T. S. Tadele, T. J. Vries, S. Stramigioli, The safety of domestic robotics: A survey of various ssafety-related publications, IEEE Robotics & Automation Magazine 21 (3) (2014) 134–142.

[21] M. A. Goodrich, A. C. Schultz, Human-robot interaction: a survey, Foundations and trends in human-computer interaction 1 (3) (2007) 203–275.

[22] M. Vasic, A. Billard, Safety issues in human-robot interactions, in: Robotics and Automation (ICRA), 2013 IEEE International Conference on, IEEE, 2013, pp. 197–204.

[23] R. Alami, A. Albu-Schaeffer, A. Bicchi, R. Bischoff, R. Chatila, A. De Luca, A. De Santis, G. Giralt, J. Guiochet, G. Hirzinger, et al., Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges, in: Proc. IROS, Vol. 6, Citeseer, 2006.

[24] V. R. Basili, G. Caldiera, H. D. Rombach, The goal question metric approach, in: Encyclopedia of Software Engineering, Wiley, 1994.

[25] P. Brereton, B. A. Kitchenham, D. Budgen, M. Turner, M. Khalil, Lessons from applying the systematic literature review process within the software engineering domain, Journal of Systems and Software 80 (4) (2007) 571–583.

[26] T. Dyba, T. Dingsoyr, G. Hanssen, Applying systematic reviews to diverse study types: An experience report, in: Empirical Software Engineering and Measurement, 2007. ESEM 2007. First International Symposium on, 2007, pp. 225–234. doi:10.1109/ESEM.2007.59.

[27] S. Keele, Guidelines for performing systematic literature reviews in software engineering, Tech. rep., Technical report, EBSE Technical Report EBSE-2007-01 (2007).

[28] C. Ghezzi, M. Jazayeri, D. Mandrioli, Fundamentals of Software Engineering, 2nd Edition, Prentice Hall PTR, Upper Saddle River, NJ, USA, 2002.

[29] H. Zhang, M. A. Babar, Systematic reviews in software engineering: An empirical investigation, Information and Software Technology 55 (7) (2013) 1341–1354.

[30] C. Wohlin, Guidelines for snowballing in systematic literature studies and a replication in software engineering, in: Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering, EASE '14, ACM, New York, NY, USA, 2014, pp. 38:1–38:10.

[31] S. Jalali, C. Wohlin, Systematic literature studies: database searches vs. backward snowballing, in: Proceedings of the ACM-IEEE international symposium on Empirical software engineering and measurement, ACM, 2012, pp. 29–38.

[32] K. Petersen, R. Feldt, S. Mujtaba, M. Mattsson, Systematic mapping studies in software engineering, in: Proceedings of the 12th International Conference on Evaluation and Assessment in Software Engineering, EASE'08, British Computer Society, Swinton, UK, UK, 2008, pp. 68–77.
URL http://dl.acm.org/citation.cfm?id=2227115.2227123

[33] D. Spencer, Card sorting: Designing usable categories, Rosenfeld Media, 2009.

[34] J. Guiochet, M. Machin, H. Waeselynck, Safety-critical advanced robots: A survey, Robotics and Autonomous Systems 94 (2017) 43–52.

[35] S. Haddadin, A. De Luca, A. Albu-Schäffer, Robot collisions: A survey on detection, isolation, and identification, IEEE Transactions on Robotics 33 (6) (2017) 1292–1312.

[36] P. A. Lasota, T. Fong, J. A. Shah, et al., A survey of methods for safe human-robot interaction, Foundations and Trends® in Robotics 5 (4) (2017) 261–349.

[37] M. Ivarsson, T. Gorschek, A method for evaluating rigor and industrial relevance of technology evaluations, Empirical Software Engineering 16 (3) (2011) 365–395.

[38] J. Popay, H. Roberts, A. Sowden, M. Petticrew, L. Arai, M. Rodgers, N. Britten, K. Roen, S. Duffy, Guidance on the conduct of narrative synthesis in systematic reviews, A product from the ESRC methods programme Version 1 (2006) b92.

[39] M. PETTICREW, L. ARAI, H. ROBERTS, N. BRITTEN, J. POPAY, Testing methodological guidance on the conduct of narrative synthesis in systematic reviews, Evaluation 15 (2009) 1.

[40] D. S. Cruzes, T. Dyb, Research synthesis in software engineering: A tertiary study , Information and Software Technology 53 (5) (2011) 440 – 455, special Section on Best Papers from {XP2010}.

[41] D. Bozhinoski, I. Crnkovic, D. Di Ruscio, I. Malavolta, P. Pelliccione, Data Extraction Form for Safety for mobile robotic systems (2016).

[42] ISO, ISO/IEC/IEEE 42010, Systems and software engineering — Architecture description (December 2011).

[43] D. C. Schmidt, Model-driven engineering, COMPUTER-IEEE COMPUTER SOCIETY- 39 (2) (2006) 25.

[44] Z. Papp, C. Brown, C. Bartels, World modeling for cooperative intelligent vehicles, in: 2008 IEEE Intelligent Vehicles Symposium, 2008, pp. 1050–1055. doi:10.1109/IVS.2008.4621272.

[45] I. GheŢa, M. Heizmann, A. Belkin, J. Beyerer, World modeling for autonomous systems, in: R. Dillmann, J. Beyerer, U. D. Hanebeck, T. Schultz (Eds.), KI 2010: Advances in Artificial Intelligence, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 176–183.

[46] L. Burgueño, M. F. Bertoa, N. Moreno, A. Vallecillo, Expressing confidence in models and in model transformations elements, in: Proc. of the ACM/IEEE 21th International Conference on Model Driven Engineering Languages and Systems 2018 (MODELS 2018). Copenhagen, Denmark, October 2018.

[47] A. Bucchiarone, A. Marconi, C. Mezzina, M. Pistore, A conceptual framework for collective adaptive systems, in: Proceedings of the 28th Annual ACM Symposium on Applied Computing, ACM, 2013, pp. 1935–1936.

[48] P. Mallozzi, P. Pelliccione, C. Menghi, Keeping intelligence under control, in: Proceedings of the SE4COG Workshop, Gothenburg, colocated with ICSE, ACM, 2018.

[49] R. Wieringa, N. Maiden, N. Mead, C. Rolland, Requirements engineering paper classification and evaluation criteria: a proposal and a discussion, Requirements Engineering 11 (1) (2006) 102–107. doi:10.1007/s00766-005-0021-6.

[50] J. C. Mankins, Technology readiness levels, White Paper, April 6.

[51] T. D. Cook, D. T. Campbell, A. Day, Quasi-experimentation: Design & analysis issues for field settings, Vol. 351, Houghton Mifflin Boston, 1979.

[52] P. Di Francesco, I. Malavolta, P. Lago, Research on architecting microservices: Trends, focus, and potential for industrial adoption, in: Software Architecture (ICSA), 2017 IEEE International Conference on, IEEE, 2017, pp. 21–30.

[53] Smart robots market - analysis & forecast to 2020, report of april 2015, Tech. rep. (2015).
URL http://www.marketsandmarkets.com/Market-Reports/smart-robots-market-48470534.html

[54] W. Afzal, H. Bruneliere, D. Di Ruscio, A. Sadovykh, S. Mazzini, E. Cariou, D. Truscan, J. Cabot, A. Gmez, J. Gorroogoitia, L. Pomante, P. Smrz, The megam@rt2 ecsel project: Megamodelling at runtime scalable model-based framework for continuous development and runtime validation of complex systems, Microprocessors and Microsystems 61 (2018) 86–95. doi:10.1016/j.micpro.2018.05.010.

[55] E. H2020, Robotics 2020 Multi-Annual Roadmap For Robotics in Europe - http://sparc-robotics.eu/wp-content/uploads/2014/05/H2020-Robotics-Multi-Annual-Roadmap-ICT-2016.pdf, 2016.

[56] F. Ciccozzi, D. Di Ruscio, I. Malavolta, P. Pelliccione, J. Tumova, Engineering the software of robotic systems, in: 2017 IEEE/ACM 39th International Conference on Software Engineering Companion (ICSE-C), 2017, pp. 507–508. doi:10.1109/ICSE-C.2017.167.

APPENDIX

Appendix A. Research Team

Five researchers carried on this study, because a 'too small' team size (e.g., single reviewer) may have difficulties in controlling potential biases [29]. Each researcher has a specific role within the team; these are identified roles:

- *Principle researcher*: PhD student with knowledge about robotics and safety aspects in software engineering; he performed the majority of activities from planning the study to reporting;
- *Secondary researcher*: an associate professor and two assistant professors with expertise in SLR methodologies, software engineering, and robotics. They were mainly involved in (i) the planning phase of the study, and (ii) supporting the principle researcher during the whole study, e.g., by reviewing the classification scheme, selected studies, extracted data, writing the final report;
- *Advisor*: senior researcher with many-years expertise in software engineering. He made final decisions on conflicts and options to 'avoid endless discussions' [29], and supported other researchers during the data analysis, findings analysis, and report writing activities.

From a geographical point of view, the research team is distributed across Belgium, Italy, The Netherlands, and Sweden.

Appendix B. Selected Primary Studies

Study	Title	Authors	Venue	Year
P1	ReFrESH: A self-adaptation framework to support fault tolerance in field mobile robots	Yanzhe Cui; Richard M. Voyles; Joshua T. Lane; Mohammad H. Mahoor	International Conference on Intelligent Robots and Systems (IROS)	2014
P2	ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots	Lynne E. Parker	IEEE Transactions on Robotics and Automation	1998
P3	Combining Quantitative and Qualitative Models with Active Observations for better Diagnoses of Autonomous Mobile Robots	Gerald Steinbauer; Franz Wotawa	Fifth Workshop on Intelligent Solutions in Embedded Systems	2007
P4	Designing autonomous robots	Saddek Bensalem; Matthieu Gallien; Félix Ingrand; Imen Kahloul; Thanh-Hung Nguyen	IEEE Robotics and Automation Magazine	2009
P5	Model-driven safety assessment of robotic systems	Nataliya Yakymets; Souhail Dhouib; Hayat Jaber; Agnes Lanusse	International Conference on Intelligent Robots and Systems (IROS)	2013
P6	An integrated model-based diagnosis and repair architecture for ROS-based robot systems	Safdar Zaman; Gerald Steinbauer ; Johannes Maurer ; Peter Lepej ; Suzana Uran	International Conference on Robotics and Automation (ICRA)	2013
P7	A versatile and safe mobility assistant	Axel Lankenau; Thomas Röfer	IEEE Robotics and Automation Magazine	2001
P8	Testing the Input Timing Robustness of Real-Time Control Software for Autonomous Systems	David Powell; Jean Arlat; Hoang Nam Chu; Félix Ingrand; Marc-Olivier Killijian	European Dependable Computing Conference	2012
P9	Building a safe care-providing robot	Leila Fotoohi; Axel Gräser	IEEE International Conference on Rehabilitation Robotics (ICORR)	2011
P10	Enhancing fault tolerance of autonomous mobile robots	Didier Crestani; Karen Godary-Dejean; Lionel Lapierre	Robotics and Autonomous Systems	2015
P11	Do whatever works: A robust approach to fault-tolerant autonomous control	David W. Payton; David Keirsey; Dan M. Kimble; Jimmy Krozel; J. Kenneth Rosenblatt	Applied Intelligence	1992
P12	Towards Rule-Based Dynamic Safety Monitoring for Mobile Robots	Sorin Adam; Morten Larsen; Kjeld Jensen; Ulrik Pagh Schultz	4th International Conference on Simulation, Modeling, and Programming for Autonomous Robots	2014
P13	SAFER: System-level Architecture for Failure Evasion in Real-time Applications	Junsung Kim ; Gaurav Bhatia ; Ragunathan Rajkumar ; Markus Jochim	Real-Time Systems Symposium (RTSS)	2012
P14	Environment rematching: Toward dependability improvement for self-adaptive applications	Chang Xu; Wenhua Yang; Xiaoxing Ma; Chun Cao	International Conference on Automated Software Engineering (ASE)	2013
P15	Formalizing correctness criteria of dynamic updates derived from specification changes	Valerio Panzica La Manna; Joel Greenyer; Carlo Ghezzi; Christian Brenner	International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)	2013
P16	Inconsistencies Evaluation Mechanisms for an Hybrid Control Architecture with Adaptive Autonomy	Bastien Durand, Karen Godary-Dejean, Lionel Lapierre, Didier Crestani	4th National Conference on Control Architectures of Robots	2009
P17	Inferring and monitoring invariants in robotic systems	Hengle Jiang, Sebastian Elbaum; Carrick Detweiler	Autonomous Robots	2016
P18	An unifying architectural methodology for robot control design with adaptive fault tolerance	W. S. Caldas; M. F. M. Campos; A. O. Fernandes; J. M. Mata	IEEE Conference on Emerging Technologies and Factory Automation	2005
P19	Handling Sensing Failures in Autonomous Mobile Robots	Robin R. Murphy; Dave Hershberger	The International Journal of Robotics Research	1999
P20	Dependable execution control for autonomous robots	Félix Ingrand; Frédéric Py	Intelligent Robots and Systems (IROS)	2004
P21	Experience with Model-Based User-Centered Risk Assessment for Service Robots	Jérémie Guiochet; Damien Martin-Guillerez ; David Powell	High-Assurance Systems Engineering (HASE)	2010
P22	ADE: A Framework for Robust Complex Robotic Architectures	James Kramer; Matthias Scheutz	Intelligent Robots and Systems (IROS)	2006
P23	Using AI techniques for fault localization in component-oriented software systems	Jörg Weber; Franz Wotawa	Mexican International Conference on Artificial Intelligence	2006
P24	Runtime monitoring of robotics software components: Increasing robustness of service robotic systems	Alex Lotz; Andreas Steck; Christian Schlegel	International Conference on Advanced Robotics (ICAR)	2011
P25	Using Controller-Synthesis Techniques to Build Property-Enforcing Layers	Karine Altisen; Aurélie Clodic; Florence Maraninchi; Eric Rutten	European Symposium on Programming	2003

P26	Timed Hazard Analysis of Self-healing Systems	Claudia Priesterjahn, Dominik Steenken, Matthias Tichy	Assurances for Self-Adaptive Systems	2013
P27	A Modeling Framework for Software Architecture Specification and Validation	Nicolas Gobillot; Charles Lesire; David Doose	4th International Conference on Simulation, Modeling, and Programming for Autonomous Robots	2014
P28	A systematic testing approach for autonomous mobile robots using domain-specific languages	Martin Proetzsch, Fabian Zimmermann, Robert Eschbach, Johannes Kloos, Karsten Berns	Advances in artificial intelligence	2010
P29	A Robot Fault-tolerance Approach Based on Fault Type	Bingu Shim; Beomho Baek; Suntae Kim; Sooyong Park	International Conference on Quality Software	2009
P30	Fault tolerant framework and techniques for component-based autonomous robot systems	Heejune Ahn; Sang Chul Ahn; Junyoung Heo; Sung Y. Shin;	Symposium On Applied Computing	2011
P31	Planning with Diversified Models for Fault-Tolerant Robots.	Benjamin Lussier; Matthieu Gallien; Jérémie Guiochet; Félix Ingrand; Marc-Olivier Killijian; David Powell	International Conference on Dependable Systems and Networks (DSN'07)	2007
P32	A methodology for testing mobile autonomous robots	Jannik Laval; Luc Fabresse; Noury Bouraqadi	Intelligent Robots and Systems (IROS)	2013
P33	Environmental Hazard Analysis - a Variant of Preliminary Hazard Analysis for Autonomous Mobile Robots	Sanja Dogramadzi; Maria Elena Giannaccini; Christopher Harper; Mohammad Sobhani; Roger Woodman; Jiyeon Choung	Journal of Intelligent and Robotic Systems	2014
P34	Rigorous System Design Flow for Autonomous Systems	Saddek Bensalem, Marius Bozga, Jacques Combaz, and Ahlem Triki	International Symposium On Leveraging Applications of Formal Methods, Verification and Validation	2014
P35	Building safer robots: Safety driven control	Roger Woodman; Alan F.T. Winfield; Chris Harper	The International Journal of Robotics Research	2012
P36	Using logic to handle conflicts between system, component, and infrastructure goals in complex robotic architectures	Paul Schermerhorn; Matthias Scheutz	International Conference on Robotics and Automation (ICRA)	2010
P37	Distributed fault diagnosis for multiple mobile robots using an agent programming language	Márcio G. Morais; Felipe R. Meneguzzi; Rafael H. Bordini; Alexandre M. Amory	International Conference on Advanced Robotics (ICAR)	2015
P38	Fault Management of Robot Software Components Based on OPRoS	JongYoung Kim; Heebyung Yoon; SungHoon Kim; Sang Hyuk Son	International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing	2011
P39	Applying the STAMP system safety engineering methodology to the design of a domestic robot	Eleftheria Mitka; Spyridon G. Mouroutsos	International Journal of Applied Systemic Studies	2015
P40	Applying regression testing to software for robot hardware interaction	Geoffrey Biggs	IEEE International Conference on Robotics and Automation (ICRA)	2010
P41	Integrating self-health awareness in autonomous systems	Karl M. Reichard	Robotics and Autonomous Systems	2004
P42	Towards declarative safety rules for perception specification architectures	Johann Thor Mogensen Ingibergsson; Ulrik Pagh Schultz; Dirk Kraft	International Workshop on Domain-specific Languages and Models for Robotic Systems	2015
P43	Towards an Ethical Robot: Internal Models, Consequences and Ethical Action Selection	Alan F. T. Winfield, Christian Blum, Wenguo Liu	Conference Towards Autonomous Robotic Systems	2014
P44	Generating certification evidence for autonomous unmanned aircraft using model checking and simulation	Matt Webster; Neil Cameron; Mike Jump; Michael Fisher;	Journal of Aerospace Information Systems	2014
P45	From AgentSpeak to C for Safety Considerations in Unmanned Aerial Vehicles	Samuel Bucheli, Daniel Kroening, Ruben Martins, and Ashutosh Natraj	Conference Towards Autonomous Robotic Systems	2015
P46	Model Based Safety Analysis for an Unmanned Aerial System	Jean-Charles Chaudemar; Eric Bensana; Christel Seguin	DRHE 2010 - Dependable Robots in Human Environments	2010
P47	Verifying Brahms Human-Robot Teamwork Models	Richard Stocker, Louise Dennis, Clare Dixon, Michael Fisher	European conference on Logics in Artificial Intelligence	2012
P48	Leveraging Collective Runtime Adaptation for UAV-based Systems	Darko Bozhinoski; Ivano Malavolta; Antonio Bucchiarone; Annapaola Marconi,	42th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)	2016
P49	Fault Tolerant Automated Task Execution in a Multi-robot System	Stanislaw Ambroszkiewicz; Waldemar Bartyna; Kamil Skarzynski; Marcin Stepniak	Proceedings of the 9th International Symposium on Intelligent Distributed Computing	2015
P50	A framework for a fault tolerant multi-robot systems	M. Tahir Khan ; M. U. Qadir ; F. Nasir ; C. W. de Silva	10th International Conference on Computer Science and Education (ICCSE)	2015

P51	Adaptive Foraging for Simulated and Real Robotic Swarms: The dynamical response threshold approach	Eduardo Castello; Tomoyuki Yamamoto; Fabio Dalla Libera; Wenguo Liu; Alan F. T. Winfield; Yutaka Nakamura; Hiroshi Ishiguro	Swarm Intelligence	2016
P52	Model-driven multi-level safety analysis of critical systems	Nataliya Yakymets ; Matthieu Perin ; Agnes Lanusse	9th Annual IEEE International Systems Conference (SysCon)	2015
P53	Online data-driven anomaly detection in autonomous robots	Eliahu Khalastchi; Meir Kalech; Gal A. Kaminka; Raz Lin	Knowledge and Information Systems	2015
P54	Improving Dependability of Industrial Transport Robots Using Model-Based Techniques	Clemens Mühlbacher, and Stephan Gspandl, and Michael Reip, and Gerald Steinbaue	IEEE International Conference on Robotics and Automation (ICRA)	2016
P55	Model-Checking and Game theory for Synthesis of Safety Rules	Mathilde Machin; Fanny Dufossé; Jérémie Guiochet; David Powell ; Matthieu Roy; Hélène Waeselynck	IEEE 16th International Symposium on High Assurance Systems Engineering	2015
P56	Towards a Virtual Machine Approach to Resilient and Safe Mobile Robots	Sorin Adam; Marco Kuhrmann; Ulrik Pagh Schultz;	IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA)	2016
P57	Hazard analysis of human-robot interactions with HAZOP-UML	Jérémie Guiochet	Safety Science	2016
P58	Measurement-based real-time analysis of robotic software architectures	Nicolas Gobillot, Fabrice Guet, David Doose, Christophe Grand, Charles Lesire, Luca Santinelli	IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)	2016

Darko Bozhinoski is a postdoctoral researcher at IRIDIA, Universite libre de Bruxelles (Brussels, Belgium). His research focuses on software engineering, software architectures, formal methods, self-adaptive systems, and robotics. As a postdoctoral researcher, he is working as part of the DEMIURGE project, a research project funded by the European Commission via an ERC Consolidator Grant. The project is focused on the design of control software for robot swarms. In this project, he is working on a formal specification language for defining swarm missions. He received a Ph.D. in Computer Science from Gran Sasso Science Institute in L'Aquila, Italy in 2017. In 2016, he was awarded the Fulbright Scholarship that enabled him to spend the academic year 2016-2017 at Carnegie Mellon University as a visiting research scholar. More information is available at http://cs.gssi.infn.it/people/bozhinoski.

Davide Di Ruscio is Assistant Professor at the Department of Information Engineering Computer Science and Mathematics of the University of L'Aquila. His main research interests are related to several aspects of Model Driven Engineering (MDE) including domain specific modelling languages, model transformations, model differencing, and model evolution. He has published more than 100 papers on such topics. Over the last decade, he has applied MDE techniques in different application domains like service-based software systems, autonomous systems, and open source software (OSS). He has been in the PC and involved in the organization of several workshops and conferences, and reviewer of many journals like IEEE Transactions on Software Engineering, Science of Computer Programming, Software and Systems Modeling, and Journal of Systems and Software. He is member of the steering committee of the International Conference on Model Transformation (ICMT), of the Software Language Engineering (SLE) conference, of the Seminar Series on Advanced Techniques & Tools for Software Evolution (SATTOSE), and of the Workshop on Modelling in Software Engineering at ICSE (MiSE). Currently, he is the technical director of the EU H2020 CROSSMINER project. More information is available at http://www.di.univaq.it/diruscio.

Ivano Malavolta is Assistant Professor at the Vrije Universiteit Amsterdam, the Netherlands. His research focuses on data-driven software engineering, software engineering for mobile development, software architecture, model-driven engineering (MDE), and robotics. He is applying empirical methods to assess practices and trends in the field of software engineering. He is program committee member and reviewer of international conferences and journals in his fields of interest. He authored more than 80 papers in international journals and peer-reviewed international conferences proceedings. He received a PhD in computer science from the University of L'Aquila in 2012. He is a member of ACM and IEEE, Amsterdam Data Science, and VERSEN. More information is available at http://www.ivanomalavolta.com.

Patrizio Pelliccione is Associate Professor at the Chalmers University of Technology and University of Gothenburg, Sweden, Department of Computer Science and Engineering. He got his PhD in 2005 at the University of L'Aquila (Italy) and from February 1, 2014 he is Docent in Software Engineering, title given by the University of Gothenburg. His research topics are mainly in software engineering, software architectures modelling and verification, autonomous systems, and formal methods. He has co-authored more than 100 publications in journals and international conferences and workshops in these topics. He has been on the program committees for several top conferences and is a reviewer for top journals in the software engineering domain. He is very active in European and National projects. In his research activity he has collaborated with several industries such as Volvo Cars, Volvo AB, Ericsson, Jeppesen, Axis communication, Thales Italia, Selex Marconi telecommunications, Siemens, Saab, TERMA, etc. More information is available at http://www.patriziopelliccione.com.

Ivica Crnkovic is a professor of software engineering at Chalmers University and Mälardalen University and a guest professor at the University of Osijek. He is also the director of Chalmers University's Information and Communication Technology Area of Advance. His research interests include component-based software engineering, software architecture, software development processes, and software engineering for large complex systems. Crnkovic received a PhD in computer science from the University of Zagreb. More information is available at http://www.ivicacrnkovic.net.