



Robotics Software Engineering: A Perspective from the Service Robotics Domain

Sergio García
University of Gothenburg | Chalmers,
Gothenburg, Sweden
sergio.garcia@gu.se

Daniel Strüber
Radboud University Nijmegen,
Nijmegen, Netherlands
d.strueber@cs.ru.nl

Davide Brugali
University of Bergamo,
Bergamo, Italy
davide.brugali@unibg.it

Thorsten Berger
University of Gothenburg | Chalmers,
Gothenburg, Sweden
thorsten.berger@gu.se

Patrizio Pelliccione
University of Gothenburg | Chalmers,
Gothenburg, Sweden
and University of L'Aquila,
L'Aquila, Italy
patrizio.pelliccione@gu.se

ABSTRACT

Robots that support humans by performing useful tasks (a.k.a., service robots) are booming worldwide. In contrast to industrial robots, the development of service robots comes with severe software engineering challenges, since they require high levels of robustness and autonomy to operate in highly heterogeneous environments. As a domain with critical safety implications, service robotics faces a need for sound software development practices. In this paper, we present the first large-scale empirical study to assess the state of the art and practice of robotics software engineering. We conducted 18 semi-structured interviews with industrial practitioners working in 15 companies from 9 different countries and a survey with 156 respondents from 26 countries from the robotics domain. Our results provide a comprehensive picture of (i) the practices applied by robotics industrial and academic practitioners, including processes, paradigms, languages, tools, frameworks, and reuse practices, (ii) the distinguishing characteristics of robotics software engineering, and (iii) recurrent challenges usually faced, together with adopted solutions. The paper concludes by discussing observations, derived hypotheses, and proposed actions for researchers and practitioners.

CCS CONCEPTS

• **Computer systems organization** → **Robotics**; • **Software and its engineering**;

KEYWORDS

robotics software engineering, interviews, online survey

ACM Reference Format:

Sergio García, Daniel Strüber, Davide Brugali, Thorsten Berger, and Patrizio Pelliccione. 2020. Robotics Software Engineering: A Perspective from the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE '20, November 8–13, 2020, Virtual Event, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7043-1/20/11...\$15.00

<https://doi.org/10.1145/3368089.3409743>

Service Robotics Domain. In *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '20)*, November 8–13, 2020, Virtual Event, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3368089.3409743>

1 INTRODUCTION

Service robots are a rising robotics domain with broad applications in many fields, such as logistics, healthcare, telepresence, maintenance, domestic tasks, education, and entertainment. A service robot is “a type of robot that performs useful tasks for humans or equipment excluding industrial automation applications” [29]. Service robots are increasingly becoming part of our lives [46], being the service robotics market estimated to reach a value of \$ 102.5 billion by 2025 [33]. Compared to industrial robotics, the service robotics domain is more challenging, since these robots usually operate in unconstrained environments, often populated by humans, requiring high degrees of robustness and autonomy.

Despite software playing an ever-increasing role in robotics, the current software engineering (SE) practices are perceived as insufficient, often leading to error-prone and hardly maintainable and evolvable software. Robotic systems are an advanced type of cyber-physical system (CPS), made up of an intricate blend of hardware, software, and environmental components. SE, despite its beneficial role in other CPS domains (e.g., automotive, aeronautics), has traditionally been considered an auxiliary concern of robotic system construction [13]. A possible reason is that robots in factory automation have built-in proprietary controllers for repetitive tasks, therefore, allowing a simple programming style. The heavy lifting is in the domains of mechanics, electronics, and automatic control [14]. In contrast, to achieve autonomy when interacting in highly heterogeneous environments, service robots are equipped with a large variety of functionalities for perception, control, planning, learning, and multimodal interaction with the human operator.

The integration, customization, and evolution of these functionalities give rise to a large amount of complexity, the management of which is a challenging task. SE systematic practices could play a crucial role in the management of such complexity, yet, according to a recent study [7], the community of SE and robotics is still not consolidated. In fact, deficiencies in the control systems of service robots stemming from a bad management of their inherent complexity are especially expensive if found after the system is put in

operation. For instance, in 2016, the Japan Aerospace Exploration Agency lost 286 million dollars due to an error in the control system of its flagship astronomical satellite. The system commanded a thruster jet to fire in the wrong direction, resulting in ten critical parts breaking off the satellite body [3].

Our long-term goal is to establish effective SE practices for a substantial lifting of the state-of-practice for engineering service robotics and provide guidance for practitioners. We also aim to help the research community to scope future research that meets industrial needs. According to Bozhinoski et al., one of the challenges “*is to promote a shift towards well-defined engineering approaches able to stimulate component supply-chains and significantly impact the robotics marketplace*” [7]. However, systematic studies about the specific software development practices and tools applied in service robotics as well as the challenges faced by practitioners in this domain are currently lacking. Academics and practitioners would benefit from such a study, as it would allow us to quantify and check existing subjective beliefs and anecdotal evidence [50], and may contribute to the formation of a body of knowledge on SE in robotics. Practitioners would further benefit from such a study as companies need to make decisions about huge investments. Having a snapshot of the current state-of-practice may help them in making informed decisions. To these ends, this paper assesses the current SE practices applied to the domain of service robotics, as well as its distinguishing characteristics and faced challenges. Our research questions are: **RQ1.** *What practices are applied in SE for service robots?* We study used SE principles, including processes and quality assurance techniques; also, we investigate technical spaces in use, including frameworks, middleware, software languages, and tools. Finally, we study reuse practices, including the frequency and motivations for reusing existing solutions or, instead, for favoring self-developed solutions. **RQ2.** *How is robotics SE perceived to be different from SE in other domains?* We identify the special characteristics that make robotics SE stand out from SE in other domains. We are interested in the effect of the robots’ features, their hardware, their requirements on autonomy, and the heterogeneity of contexts where they are deployed. **RQ3.** *What challenges do practitioners face related to SE for service robots?* We elicit the most common SE challenges faced by practitioners in service robotics, as well as the solutions applied.

To collect data, we conducted 18 semi-structured interviews with industrial robotics experts working for 15 companies from 9 different countries. We accompany this study with an online survey, targeting industrial and academic practitioners in the robotics domain, from which we collect 156 responses. Academic practitioners are those who work for an academic institution (we recruited most of our respondents from significant GitHub repositories with robotics-specific code). Among the survey respondents, we identified practitioners working for 58 companies from 20 different countries and 16 academic institutions from 10 different countries. To the best of our knowledge, our study is the first with this ambition.

In our study, we discovered that roboticists are predominantly focused on implementation and real-world testing, often favored over simulation. We learned that robotic control systems are typically developed as component-based systems, implemented by developers who may come from different backgrounds (e.g., mechanical, electrical, or software engineers). We also elicited the main characteristics of robotics SE, where the cyber-physical nature of robots

and the variety of disciplines required to develop a complete robotic system were highlighted. These characteristics increase the complexity of robots’ control software, calling for systematic practices as modeling and the usage of software architectures to improve the development process [18]. Our respondents ranked challenges related to robustness and validation as most pressing, and typically address them by applying thorough testing processes. Based on our observations, we also identify research themes that deserve further investigation. We provide conjectures for why these themes are currently under-investigated and recommendations for both researchers and practitioners. In summary, we contribute:

- The findings from our interview data, including qualitative information about applied practices, distinguishing characteristics, and faced challenges and solutions in service robotics.
- The results from our online survey, which supports a quantitative assessment of the trends regarding SE practices in robotics.
- A replication package [25] containing (1) the questionnaire and the anonymized responses from our online survey (2) our interview guide and used open-ended questions, (3) the codebook from the qualitative data analysis, (4) and our quantitative data analysis scripts.
- Discussion of observations, derived hypotheses, and proposed actions for researchers and practitioners.

2 RELATED WORK

Current achievements in defining standardized solutions, reference architectures, and model-driven approaches for the development of reusable robotic systems are documented in various survey papers briefly summarized in the following.

Brugali et al. [10] present five challenges that characterize robotics, namely (1) the identification of stable requirements of robotic systems, (2) the definition of abstract models to cope with hardware and software heterogeneity, (3) the seamless transition from prototype testing and debugging to real system implementation, (4) the deployment of robotic applications in real-world environments, and (5) the design and integration of software control systems depending on highly heterogeneous technologies. Brugali et al. [8, 9, 11] have reviewed existing approaches that exploit component-based SE (CBSE) and Model-Driven Engineering (MDE) in robotics. Notable examples of component models and associated MDE toolchains specifically defined for developing robotic applications are OROCOS [16], SmartSoft [42], BRICS [15], and LAAS-BIP [5].

MDE has also been applied recently in service robotics as a core technology to compose systems out of exchangeable components and move towards well-engineered system development processes as opposed to current craftsmanship practices [46]. Steck et al. [47] propose a system that exploits design- and run-time models to build robotic systems and support them in their decision-making process. Closely related to MDE are Domain-Specific Languages (DSLs), which have also been highlighted for their benefits, including the improvement of composability and system integration [46]. The research community has already worked in the last years on the use of DSLs for the development of robotic systems [15, 43, 44, 47]. Nordmann et al. [36] present a survey of more than 130 papers on DSLs in robotics. The papers are classified according to robotic functional subdomains, such as Kinematics and Motion Control.

Elkady et al. [23] present a literature survey and attribute-based bibliography of the state-of-the-art in robotic middleware frameworks. For example, the Middleware for Mobile Robot Applications (MIRO) [51] provides a robotic hardware abstraction layer that simplifies the development of robot control applications.

Ahmad et al. [1] conducted a systematic mapping study on software architectures for robotic systems, which reveals that the conventional concepts of architectural frameworks and architectural notations are well integrated into robotics. However, they also identified a lack of focus on architecture-specific evaluation for validating requirements for robotic software. In addition, Kortenkamp et al. [32] present various approaches to architecting robotic systems with an emphasis on programming tools and environments.

Ingrand et al. [28] provide a survey that introduces a global perspective of autonomous robots' deliberation—i.e., endowing robotic systems with adaptable and robust functionalities.

Estefó et al. [24] interviewed and surveyed Robot Operating System (ROS) [41] developers. They report on the issues related to resources reuse within this ecosystem.

Bozhinoski et al. [7] provide a systematic mapping study on the state of the art in safety for mobile robotic systems. Focusing on an SE perspective, they classify safety solutions from 58 primary studies concerning trends, characteristics, and industrial adoption. They found that existing solutions are largely not ready to be used in uncontrollable and unknown, often shared with humans environments. The only source of information of this study consists of published papers that have been selected as primary studies. This motivates us to perform a study to better investigate these aspects and by involving industrial and academic practitioners in robotics.

Synergies with SE have been investigated also for other domains that share common characteristics with robotics. Amershi et al. [4] conducted a study observing software teams at Microsoft as they develop AI-based applications. Altinger's Doctoral Thesis [2] reports on fault prediction techniques applied to automotive software development. Bures et al. [17] report on challenges and promising solutions in the area of SE for Smart Cyber-Physical Systems.

3 METHODOLOGY

Figure 1 summarizes our methodology. For an initial set of four exploratory interviews, we relied on our own experiences as SE researchers in various robotics projects [21, 22] where we collaborated with robotics engineers, many of them from industry. In those interviews, we asked previous collaborators, who are expert robotics engineers, about their software engineering practices, tools, and challenges, following the typical software-engineering phases (from requirements via design over implementation and quality

assurance), while also aiming to understand the organizational and business context. This helped us sharpening our knowledge of the typical terminology used in robotics. Based on this knowledge and our research questions, we designed an interview guide and, in parallel, a survey questionnaire for the second phase of our study. Our replication package [25] provides the questionnaire, an analysis script, detailed results, and our interview guide.

3.1 Survey

Design. We implemented the survey questionnaire with Google Forms, designing it to take no more than 15–20 minutes to complete. It consists of 47 questions (27 optional, 20 mandatory) divided into the sections: robotic platforms and context (e.g., which types of robots are engineered), activities and paradigms (e.g., to which extent requirements engineering is performed or what tools are used), development practices (e.g., whether and how software artifacts are reused), quality assurance (e.g., where test data is obtained from), mission specification (e.g., whether the robot behavior is hard-coded), challenges and solutions (e.g., concerning applying AI techniques), general perception (e.g., what characteristics are unique to the domain), and finally a section on the respondents' demographics (e.g., years of experience). The majority of questions was closed-ended, using either checkboxes or 5-point Likert scales, the latter to express the frequencies of applying certain activities (e.g., code review), ranging from never via almost never, sometimes and very often, to always. We also allowed declaring "Don't know."

With this design, we followed an inductive and a deductive strategy: on the one hand (inductive), we used questions similar to the ones asked during the interviews so to generalize our findings; on the other hand (deductive), we asked about hypotheses (proposed explanations for phenomena) we learned about in the exploratory interviews. Some examples of such hypothesis are: (1) robotics software developers tend to "reinvent the wheel" instead of reusing already existing solutions; or (2) some of the most important challenges in the field are the lack of standardized solutions and of internal and external documentation. To this end, we designed questions such as: "Have you ever developed a software component from scratch rather than reusing an existing (either self-developed or third-party) one? If yes, why?" After the initial design, we tested the questionnaire internally among the authors, and five times externally with industrial and academic colleagues with expertise in robotics to fine-tune length, phrasing, and terminology before distribution.

Distribution. Our distribution strategy was three-fold. First, we used social media (Twitter and LinkedIn) to distribute our survey invitation, clearly asking for responses from practitioners in robotics. Social media stimulated the response rate, since it is inherently interactive and practitioners are active there—e.g., we knew examples from Twitter and LinkedIn allowed searching for specific companies and roles. Second, we created a list of contacts from our previous projects and collaborations. Third, we mined GitHub repositories to extract their contributors' public email addresses. Specifically, we sought repositories from well-known tools (e.g., Gazebo [31]), frameworks (e.g., ROS [41], ROS2 [34], OROCOS [16], Yarp [35]), and libraries (e.g., OpenCV [40]). We also looked into repositories of well-known companies (e.g., Boston Dynamics), companies found in the media [19], and companies from survey respondents who

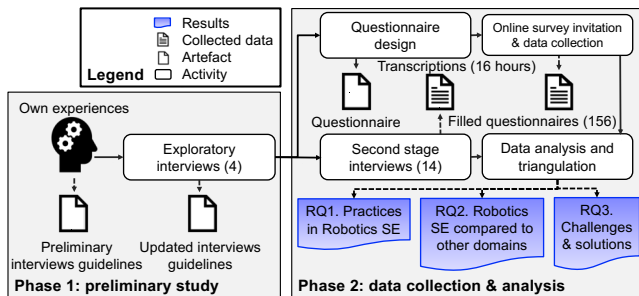


Figure 1: Research methodology overview

provided their email addresses. We also asked recipients to distribute our survey to other potential respondents. Finally, to improve gender balance, we used Google to find female robotics engineers and found articles on this topic [49].

Survey Respondents. We sent out 2499 invitations and collected 156 responses (response rate of 6%, in line with the expectations of SE survey’s response rate [45]). The most common occupation among the respondents is industrial practitioner with 60.9% (31.4% of leading technical roles and 29.5% of programmers). We also received a substantial number of responses from academic practitioners (57, 36.5% of the total), so we decided to not filter them out and instead use that data to compare practices applied by academic and industrial practitioners. The last group is formed by practitioners who did not identify themselves as industrial nor academic, and people who decided not to answer with their occupation (4, 2.6%).

We used provided email addresses to identify the companies for which our respondents work, and the associated countries. That way, we mapped 72 of our industrial practitioners to 58 companies from 20 different countries, distributed over geographical areas as follows (number of respondents between parenthesis): a) Europe (51), b) North America (13), c) Asia (7), and d) South America (1). We also mapped 22 of our academic practitioners to 16 universities from 10 different countries from the following geographical regions: a) Europe (18), b) Asia (1), c) South America (1), and d) Oceania (1).

Almost half of our respondents have between five and 10 years of experience in robotics (43.6%), 17.3% between 10 and 15, and 11.5% more than 15. Out of the total, 24.4% of the respondents have between one and five years of experience and 3.2% did not provide an answer to this question. Regarding the size of the company, almost half of the answers (46.8%) belong to respondents who work for companies with 50 or more employees. Then, 11.7%, 22.1%, and 15.6% of the respondents work for companies with 31-50, 11-30, and 1-10 employees, respectively. Among all the respondents, 91.2% are men, 6.1% are women, and 2.7% preferred not to specify. In the remainder, we refer to individual survey respondents as S1, S2, and so on.

3.2 Interviews

Design and Conduction. We designed an interview guide targeting interviews of around 45-60 minutes length. The guide started with questions regarding the interviewee’s background (e.g., subject of university studies and previous roles at companies) and roughly resembled the questionnaire’s structure and addressed each research question. The interviews were semi-structured: we had question blocks we wanted to cover, but allowed the interview to flow freely based on interesting aspects the interviewee brought up, where we then dug deeper if necessary. Before each interview, we also studied the survey responses to prepare specific questions, such as for clarification. We also emphasized special topics depending on the interviewee and their role at the company. For instance, CEOs were usually able to provide much information regarding the company’s organization or the overall project life-cycle. Software and robotics engineers mostly provided information on development issues, engineering paradigms, and technological spaces. Following a grounded-theory-like approach for the analysis (explained shortly), we also slightly changed our attention to frequent topics that arose along with the interviews, such as the dominant use of ROS as a robotics framework, its causes, and benefits and shortcomings.

Table 1: List of interviewees

	Cp. ¹	Exp. ²	Role ³	Company description
1	A	6	RSE	Small, specialized in outdoor mobile robots and professional Research and Development (R&D) consultancy.
2	B	18	CEO	Small, focuses on autonomous navigation solutions.
3	C	5	PM	Large, mainly works in the field of engineering and electronics. Currently investing in artificial intelligence and robotics.
4	C	6	PM	
5	D	15	RSE	Large technological center, mainly works in applied R&D.
6	E	2	RSE	Small, manufactures humanoid robots and robotic platforms.
7	F	8	RSE	Small, manufactures mobile service robotic platforms.
8	G	5	RSE	Large, works in the intralogistics sector.
9	H	2.5	RSE	Technology-oriented large company, also invests in robotics.
10	I	6	RSE	Large technology institute, which focuses in diverse industrial fields, including mobile robots.
11	C	14	RSE	
12	J	8	SAE	Medium-sized, provides self-driving products and services.
13	K	4	CEO	Micro-business, develops and rents customer-tailored robots.
14	L	4	CEO	Micro-business, develops and rents modular robot solutions.
15	M	21	PM	Large, mainly works in the aviation sector, but also with unmanned air vehicles.
16	N	11	PM	Large research institute, focuses on industrial projects like artificial intelligence in automation.
17	J	7	RSE	
18	O	8	CTO	Small, works in educational technology using service robots.

¹Company name (letter code), ²experience with robotics in years.

³Roles: Robotics Software Engineer (RSE); Chief Executive Officer (CEO); Project Manager (PM); Application Engineer (AE); Chief Technology Officer (CTO).

Interviewees. We recruited our interviewees: (1) among the survey respondents who agreed to be contacted for follow-up questions, (2) targeting domain experts using social media (see explanation in Sec. 3.1), and (3) among previous colleagues with expert robotics knowledge. We performed 18 interviews, summing up to 16 hours of interview data. As detailed in Table 1, the interviewees belong to 15 companies from nine different countries (16 interviewees work in Europe and two of them in North America), with roles ranging from developer to CEO, giving us a broad perspective. We made sure that all interviewees are pure industrial practitioners. The table also shows our interviewees’ overall working experience within robotics and briefly characterizes each company. In the remainder of the paper, we refer to individual interviewees with I1, I2, and so on.

3.3 Data Analysis

We analyzed the survey quantitatively by creating diagrams and manually aggregating responses to questions. To understand trends in the Likert-scale questions, we created bar charts (many of which are shown in the remainder). Furthermore, we analyzed the survey results qualitatively by inspecting responses to the open-ended questions. We created a technical report with the detailed results (available in our replication package [25]), also separating the results by the respondent’s role, specifically: industrial practitioner (further sub-divided into leading technical role and programmer) and academic practitioner. We make this separation to understand any deviations based on the role, and if there are, we report relevant differences separately. After creating the technical report, we triangulated the results with those from the interviews as follows.

We analyzed the interviews by transcribing them and then performing open coding [48], as known from grounded theory [20]. Two authors coded the interviews iteratively, continuously discussing and refining the codes and organizing them in a hierarchy. We defined a codebook consisting of 373 codes, which we provide in

our replication package [25]. Based on the codes, we triangulated results from the survey and interviews (mapping codes to parts of the technical report about the survey data). We report the results by first providing the context in which our survey respondents and interviewees perform their engineering, shortly in Sec. 3.4. Thereafter, in Sections 4, 5, and 6, we use a narrative style to answer our research questions RQ1, RQ2, and RQ3, respectively.

3.4 Robotic Platforms and Context

Most respondents (75%) affirm that they use (ground) mobile robots in their projects. Collaborative robots and mobile manipulators come in second and third place with 45.5% and 42.3%, respectively.

Table 2 gives an overview of the types of products and services provided by the respondents and what field they are applied to. We find a predominance of planning modules (64.1%) and low-level functionalities (63.5%) as the most common types. Industrial practitioners tend to provide more complete robotic systems (59.6%) than the academic ones (37.5%). Our survey allowed us to enter an open-ended text after answering “Other,” but despite the number of people who chose this option only 13 open-ended responses were collected for this question. Examples of answers include testing systems, real-time operating systems, and high-level vision solutions. The most common application fields are factory automation and research for service robots (48.1% for both). The data shows a tendency among industrial practitioners towards factory automation (51.1%) while academic practitioners favor research for service robotics (60.7%).

4 PRACTICES (RQ1)

RQ1 focuses on practices applied during the software development life cycle of robotic projects. It includes typical activities, development paradigms, development processes, software languages and frameworks, quality assurance (QA), reuse practices, and used tools.

4.1 Activities

Figure 2 (a) gives an overview of the frequency of different activities according to the survey. The most common one is implementation, 99% of the respondents perform it always or very often. Real-world experimentation and testing and simulation come in second (88%) and third place (79%). The least common activities are automatic code generation (15%) and requirements engineering (RE) with 49%.

Real-world experimentation and testing and simulation are performed equally by industrial and academic practitioners. Interestingly, respondents perform real-world experimentation more often than simulation (88% versus 79%). The difference between industrial and academic practitioners for software maintenance and evolution

Table 2: Provided service or product and application fields

	Overall	Industrial pract.	Academic pract.
Planning modules	64.1 %	58.5 %	75 %
Low-level functionalities	63.5 %	59.6 %	71.4 %
Complete system	51.3 %	59.6 %	37.5 %
Drivers	44.9 %	43.6 %	44.6 %
Other	9 %	9.6 %	8.9 %
Factory automation	48.1 %	51.1 %	42.9 %
Service robotics research	48.1 %	38.3 %	60.7 %
Others	34.6 %	40.4 %	26.8 %
Transportation	24.4 %	24.5 %	23.2 %

mounts to a rate of 73% and 56%, respectively. The gap may be related to work conventions in academia: academic practitioners do not focus on providing a final product to a customer, but on generating a proof-of-concept design that can be used to validate a scientific contribution. Only 8% of the respondents responded “Don’t know” to the RE question and, for the reminder, the respondents that perform RE are surprisingly low. We believe that the respondents had a very specific understanding of RE in the sense of using dedicated RE tools. Another possible reason is the large number of participants who follow agile processes. We analyzed our data and discovered that, from 156 respondents, 70 perform RE, and 48 perform RE *and* agile processes always or very often. Therefore, at least a significant portion of respondents who use agile processes indicate RE experience. Since they all need to somehow elicit requirements, we dug deeper in the interviews and in fact we found that the majority elicit requirements but in rather ad hoc ways. I8: “*We had the customer requirements and also all the list of features the customer needed, and then we again sat down and decided how to design the software architecture and realize that. But there was no official process that we followed.*” Furthermore, even when they followed an RE process, some stated they did not, probably due to a terminology issue.

4.2 Development Processes

We present an overview of the frequency of applied SE processes in Fig. 2 (b). The most common development process applied by the survey respondents is agile (73% of the respondents apply it always or very often versus 21% and 17% for waterfall and hybrid, respectively). The “Don’t know” answers’ rate is 11%, 26%, and 30% for agile, waterfall, and hybrid, respectively. There is a remarkable difference between the industrial and academic practitioners’ results: 81% of the industrial ones adopt agile processes very often or always, while the percentage of academic ones is 60%.

We observe that agile SE processes are more often followed by industrial practitioners than by their academic counterparts (81% against 60%). Three academic practitioners respondents even question how systematic are the processes followed by academia. S1: “*Almost always there is no clear SE process.*”

Practitioners mainly focus on implementation and real-world experimentation (preferred over simulation) during software development, typically following agile-style processes.

— Observation 1 (Activities and Development Processes)

4.3 Paradigms

Figure 2 (c) gives an overview of the SE paradigms and their use frequency in robotics. Most survey respondents make use of object-oriented programming (OOP) (92%) and many follow a CBSE approach (65%). Software product line engineering (SPLE) and model-based software development (MBSD) are rarely used (12% and 25% use them always or very often, respectively). The “Don’t know” answer’s ratio is 32% and 15% for SPLE and MBSD, respectively.

The prevalence of OOP and CBSE can be explained with the wide usage of ROS, which works with Python and C++ and enforces a component-based approach. We observe that only industrial respondents report frequent use of SPLE, specifically 19% of the industrial practitioners report using it very often or often, versus 0% of academic practitioners. We also conclude that MBSD is not extensively

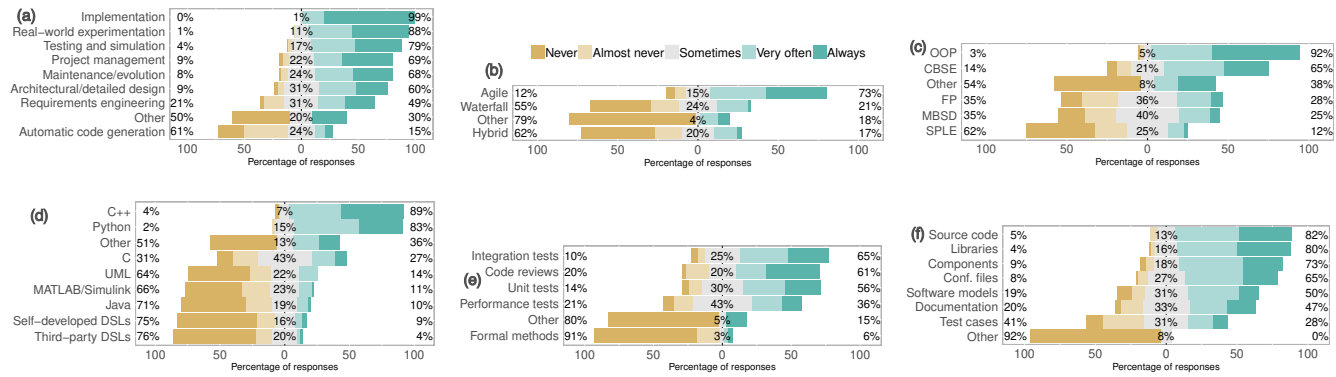


Figure 2: Survey Results. Reported Frequencies of (a) Activities; (b) Processes; (c) Paradigms; (d) Software Languages; (e) QA Techniques; and (f) Reuse of Artifacts. We report on cases with a significant share of “Don’t know” answers in the text.

used in robotics, especially among leading technical role practitioners (only 16% of them apply it very often or always, in contrast to 30% and 27% of industrial programmers and academic practitioners).

4.4 Software Languages and Frameworks

Figure 2 (d) provides an overview of the use of software languages in robotics. C++ and Python are used most often; 89% and 83% of our survey respondents use them very often or always, respectively. Self-developed (9%) and third-party (4%) domain-specific languages are the least frequently used among the options. The “Don’t know” answer’s rate for C++ and Python is 4%, while it ranges from 14% to 23% for the rest of the options.

Considering frameworks, ROS is undeniably the most used framework for robotics software development nowadays. In particular, 88.5% of the survey’s respondents and 16 interviewees use it. ROS2 and OROCOS follow as second and third with 22.4% and 18.6%, respectively. ROS2 is, according to the conducted interviews, not yet implemented in most of the companies, and there is not much initiative to change that. An important reason that drives that decision is stated by I16: “Some of our industrial partners are also more conservative. Never change a running system. If you have now ROS1, [...] why change it if it’s not clear what exactly would be the improvement?”

For either question, there are no big differences in the results between academic and industrial practitioners, probably due to the widespread use of ROS in both communities. ROS is mainly written in C++ and Python and provides many libraries that use them. According to one interviewee, Python is normally used for high-level components that do not have hard-time constraints or for prototyping, since it is simpler than C++. I7: “It’s really fast and dynamic. [...] For better performance we use C++, and also with C++ we have the possibility to better protect the code than with Python.” However, a considerable number of the responses (around 20%) denote a lack of awareness with respect to the other used languages.

Practitioners usually rely on frameworks to develop software, and ROS is the preferred framework among roboticists. Aligned with the popularity of ROS, we find C++ and Python as the most used software languages, and OOP and CBSE as most typical paradigms.

— Observation 2 (Paradigms, Languages, and Frameworks) —

4.5 Quality Assurance

We investigate quality assurance (QA) techniques and sources of test data for industrial and academic practitioners.

Techniques. Figure 2 (e) shows that the most common QA technique among respondents is integration testing, 65% of them apply it very often or always. Code reviews is the second option (61%), unit testing the third (56%), and performance testing is the fourth (36%). The least common answer is formal methods (6%) and in fact 43% answered “Don’t know” to this question.

Integration testing may be the most applied QA technique (65%) due to the popularity of ROS, which enforces a component-based approach for developing software. Therefore, practitioners need to integrate different components to test a final product. Oppositely, industrial practitioners often work in teams. Surprisingly, performance testing is not performed as often as we expected taking into account the cyber-physical nature of robots and the performance constraints (e.g., regarding memory, power consumption, computation requirements) they entail.

The most widespread QA techniques are code reviews and testing, especially integration testing. Surprisingly, despite the role of hardware constraints, performance testing is rarely performed.

— Observation 3 (Quality Assurance Techniques) —

Sources of Testing Data. The survey’s results show that the most common source of data for testing is simulation (82.1%). The second most used source is runtime monitoring (75%) and the third manually based on experience (64.7%). Testing data is acquired from runtime monitoring (75%) almost as often as from simulation (82.1%), which may mean that both industrial and academic practitioners perform almost as much simulation as real-world testing. This was already identified in Fig. 2 (a) and will be further discussed in Sec. 6.

Practitioners’ main source of data is simulation, but the difference with runtime monitoring is surprisingly low.

— Observation 4 (Sources of Testing Data) —

4.6 Reuse

Reused Artifacts. As shown in Fig. 2 (f), source code and libraries are the most commonly reused artifacts: 82% and 80% of the respondents stated that they always or very often reuse them. Software components and configuration files come in third and fourth with 73% and 65%. The percentages are similar for all the roles concerning these artifacts. On the other hand, the least reused type of artifact is test cases, which are always or very often reused by 28% of the respondents. Among the industrial programmers, 38% very

often or always reuse test cases versus 21% and 23% of leading roles and academic practitioners, respectively.

Very few respondents reuse documentation (47% of the respondents reuse it always or very often), in relation to the reuse frequency of source code, libraries, or software components (82%, 80%, and 73%, respectively). This finding relates to a challenge we discuss in Sec. 6: lack of documentation and how it hinders further software reuse (obs. 15). On the other hand, industrial programmers are the group that most commonly reuses documentation (51%, 8% higher than academic practitioners, and 9% higher than leading roles) and test cases (38% as opposed to the 28% of the overall). A possible explanation of these findings is that programmers are primarily concerned with understanding, writing, and testing of code.

Roboticians commonly reuse source code and libraries, but not so often test cases and documentation.

— Observation 5 (Reused artifacts)

Roadblocks for reuse. The main reason (78.2%) among the survey respondents to not reuse existing software resources relates to technical problems (e.g., component’s granularity does not fit, missing functionality, incompatible interfaces). The second and third most common reasons are lack of documentation (46.8%) and licensing issues (44.2%). Among the respondents, 19.9% favor self-developed components. Data regarding technical problems and licensing issues differ between industrial and academic practitioners: the first group finds more technical problems and licensing issues than the second (83% versus 69.6% and 53.2% versus 28.6%, respectively).

The second most common cause for not reusing existing software components is lack of documentation (46.8%), which might be related to the first cause (technical problems with 78.2%)—i.e., it is hard to assess the specifications of a component without proper documentation. These findings are in line with literature: According to a recent study [24], technical problems are the most common issue (e.g., outdated and buggy resources) for reusing software in ROS—often caused by a lack of maintenance. The same study points the lack of documentation as a problem for 56% of respondents. As expected, licensing issues are more problematic for industrial practitioners than for academic practitioners (53.2% versus 28.6%).

The main roadblocks for reusing software artifacts are (mainly) technical problems and lack of documentation.

— Observation 6 (Roadblocks)

Interoperability. Remarkably, for interoperability among software components of their robotic systems, practitioners tend to rely

Table 3: Most used tools for: Implementation (I), Real-World Experimentation (RWE), Testing/Simulation (TS), Project Management (PM), Software Maintenance/Evolution (SME), Architectural and Detailed Design (ADD), Requirements Engineering (RE), Automatic Code Generation (ACG).

I	Visual Studio (22%)	QtCreator (10%)	Vim (8%)	Pycharm (7%)
RWE	Roslaunch (24%)	Rosbag (20%)	In-house (11%)	ROS (8%)
TS	Gazebo (36%)	In-house (12%)	gtest (8%)	Rviz (3%)
PM	Jira (19%)	Git (18%)	Trello (14%)	Excel (11%)
SME	Visual Studio (20%)	Git (18%)	Eclipse (7%)	Vim (6%)
ADD	UML editors (20%)	PowerPoint (13%)	Draw.io (10%)	Confluence (5%)
RE	Word (18%)	Excel (17%)	Jira (14%)	Confluence (13%)
ACG	In-house (19%)	Python (12%)	Matlab (8%)	None (8%)

on the experience of the team (72.4% of the responses). This answer was selected by far more often than those related to a systematic approach to interoperability, such as having a clear definition of the interfaces of each component (42.9%), following a precise architecture (26.3%), or following a standardized solution (10.9%). The most common standardized solution stated by the survey respondents is ROS itself and its set of standardized messages.

To ensure the interoperability among heterogeneous system components, practitioners tend to rely on ad hoc practices.

— Observation 7 (Interoperability)

4.7 Tools

Table 3 shows the four most used tools among the respondents for each activity introduced in Fig. 2. The use of dedicated tools for project management indicates a high level of commitment to this activity. Contrary to industrial practitioners, the academic ones favor open-source tools as Redmine (9.9%) against proprietary ones as Jira (6.9%). Respondents (both academic and industrial practitioners) use mostly office suite applications for RE. This indicates that requirements are usually elicited, specified, and modeled in a simple manner, without using dedicated software. UML editors (e.g., Papyrus, TikZ-UML, PlantUML), PowerPoint and draw.io are the most common tools for Architectural and Detailed Design.

The preference of VisualStudio and PyCharm for implementation can be related to the preference of C++ and Python, the two main languages supported by the ROS framework. Gazebo is the most popular tool for testing and simulation (36%), which is not surprising considering its strong integration with ROS. The answers related to the software maintenance and evolution activity are similar to those seen for implementation.

Respondents use specialized tools for the most common activities (e.g., Gazebo for simulation) and generic tools for less common ones (e.g., the Office suite for requirements engineering).

— Observation 8 (Tools usage)

5 DISTINCT CHARACTERISTICS (RQ2)

RQ2 focuses on the distinguishing characteristics of robotics SE compared to other domains, particularly cyber-physical domains. These are characterized by systems that integrate physical environment and hardware with software—e.g., automotive, avionics. We discuss the responses to an open-ended question, in which we asked the survey respondents to name and explain such characteristics. We analyzed the respondents’ answers to create codes, group all answers (and relevant parts thereof) using the codes, and counted the matches of answers to codes. Overall, 156 textual answers were provided and 161 matches counted. In the following summary, we triangulate the survey results with answers from our interviewees. **Cyber-Physical Systems.** A total of 64 of our survey respondents perceive robotic systems as cyber-physical systems and stress the integration between physical aspects and hardware and software components. S2: *“Any small change in code can be seen easily in the behavior of the robot. On top of that, the work is done really close to the hardware that constitutes the robots.”* Interactions with the real world lead to a special emphasis on non-functional concerns such as safety and reliability: I14: *“I think we have to be very careful when we develop software for systems that are out there in the physical*

world. If I don't get a particular tweet, that doesn't really matter much. But if my robot goes berserk and starts killing people, that's not good."

Heterogeneous Disciplines. Out of the total, 24 survey respondents noted that a robotic system bridges diverse domains such as hardware drivers, firmware, planning, perception, and AI. S3: "The systems we develop are incredibly complex and involve multiple disciplines and technologies, while doing work in the real world. This means there are dozens of integration points and world interactions that have to be designed, engineered, and tested." Our synthesis of the topic is:

The creation of a full-working robotic system requires knowledge beyond just software, including hardware, mechanics, and physics.

— Observation 9 (Heterogeneous Disciplines) —

On the organizational level, the associated challenges are addressed by engineering teams with diverse experiences and backgrounds.

Testing. As observed by 22 respondents, the interaction with the physical world gives rise to specific requirements on the testing process. They stress a need for integration testing in the real world. S5: "Testing is slower. [It] Requires that you actually have a piece of hardware operating in the real world, where every sensor is lying to you, your environment is actively fighting back, and you have to still maintain sufficient functionality for an outside observer to be happy with your performance. Every assumption must be constantly questioned".

Maturity. A total of 20 survey respondents address the maturity of the robotics SE community. Awareness for SE practices seems to be lower than in other domains. S6: "Common SE practices are less common or known in robotics and therefore software and software maintainability suffer." Our observation is as follows:

Awareness of software engineering practices among robotics practitioners is low, which hinders software quality and maintainability.

— Observation 10 (SE practices) —

A possible explanation is that software development is often seen as an auxiliary task for developing an actual system. S7: "Software is not the goal at all. It's only an executable form of robotics theories and algorithms." A recurring topic mentioned by three survey respondents is a lack of standardized solutions. S8: "In robotics, SE is challenged by the lack of standards and guidelines in the development."

Context and Adaptation. Real-world interactions lead robots to deal with heterogeneous contexts and adaptation, topic that is touched on by 13 survey respondents. S10: "Robots act in the real world, which introduces all sorts of unexpected error cases that one wants to catch. Software engineering in other domains often feels «sandboxy» in comparison." Various responses point out that a robot also becomes an actor that shapes the environment. S11: "They also

change the environment itself via actions. Covering all these cases might be intractable in certain scenarios."

Among the most frequently named characteristics of robotics are (1) their cyber-physical nature, (2) the need for testing processes that contemplate the influence from hardware and the physical world, (3) and the necessity to react to a large variety of contexts.

— Observation 11 (Adaptable Cyber-Physical Systems) —

No difference. In contrast to the majority of responses, only eight survey respondents found the differences to other domains negligible. One interviewee points out that an emphasis on the differences to other domains can actually lead to a problematic mindset. I5: "It differs much less than people think. [...] Actually, sometimes robotics reinvent the wheel or have a syndrome of «not invented here»."

Influence of academia. The influence of academia is commented on by four survey respondents, mostly pointing out a negative influence of academia: S12: "A larger than usual share of code comes from researchers, as opposed to professional software developers. This often causes quality issues."

With respect to other CPS. What follows are the main differences between robotics and other CPS (e.g., automotive) as perceived by four responses: (1) Higher autonomy, more decision making and planning; (2) More flexibility for interacting with the world; (3) More generic perception capabilities; (4) More advanced and computationally intensive algorithms (for perception and planning).

The main differences with respect to other CPS are a higher importance of computation requirements and autonomy for decision making, since robots interact with the environment and humans.

— Observation 12 (Distinguishing Characteristics related to other CPS) —

6 CHALLENGES AND SOLUTIONS (RQ3)

Practitioners in service robotics face various challenges and apply a large variety of solutions to address them. In this section, we first present in detail nine of the major challenges that we identified during the first phase of our study (see Fig. 1) from the exploratory interviews, our experiences, and literature (specifically, the studies of Brugali [10], Bures [17], and García [26]). We quantified the relevance of these challenges in the study's second phase through the survey. We report on a rating of the challenges relevance's quantification (see Fig. 3 for an overview), and on their applied solutions as specified in open-ended questions. Coding 143 responses to the open-ended question yielded 150 solutions; we filtered out 13 answers for being empty or due to not specifying a solution.

Robustness. Most respondents (83%) agreed or strongly agreed that the biggest challenge is ensuring enough robustness to robotic systems. That is, both software and hardware parts that constitute a robot must work continuously without (critical) failures regardless the context—i.e., the mission and the environment. I3: "Programming the [robot's] task is not the most time consuming thing, but programming the failure handling [...] People might need two weeks to set up one application, two days to set up the task, and eight days to set up all the failure case." We collected 19 solutions for ensuring robotic robustness. The most frequently named ones are rigorously defined development processes (e.g., V-Model) and thorough testing.

Validating Robotic Systems. It appears to be the second most important challenge, as acknowledged by 74% of our respondents.




Figure 3: Challenges in robotics software engineering.

Validation is challenging due to the interaction of robotic systems with the real world, which makes the validation process both expensive and slow (obs. 11). We obtained 16 responses with solutions, including the employment of an independent team for testing, defining a standard set of tests the robot or middleware has to be compliant to, and formalizing (e.g., modeling) the system as much as possible to automate code generation and validation processes.

Lack of Documentation. It was considered a major challenge by 63% of our respondents. This finding is aligned with our observation from Sec. 4 (obs. 6): 46.8% of the respondents do not reuse already existing solutions due to lack of documentation. I1: “The biggest challenge for a commercial company is actually to get the right documentation or the right engineers having the right experience and knowledge about those tools.” In 17 responses with solutions, we find the reviews, use of wikis, guidelines on comments, and automatic generation of documentation as typical solutions. Five respondents consider this challenge as unsolved, as developers frequently have to understand undocumented source code.

Dynamic Adaptation. It is considered a challenge by 62% of the respondents. Dynamic adaptation is becoming an important topic for academic and industrial practitioners, as service robots are expected to adapt to environment conditions [26]. This challenge is tightly related with robustness (i.e., autonomy). From nine solutions specified by our respondents, the most common one (three responses) is to test as many real-world situations as possible, focusing on possible problems and outcomes, which are then modeled and hard-coded in the system to enable adaptation. A common modeling paradigm (mentioned by four interviewees) is to define adaptation based on states and pre-defined events, as in state machines or behavior trees.

Achieving robustness, validation, and dynamic adaptation are ranked as three of the top 4 challenges. The main solutions for these challenges are rigorous development and test processes (involving real-world testing) and programming error handling.

— Observation 13 (Robustness and validation) —

Interoperability. This topic is considered a challenge by 61% of the respondents. As discussed in Sec. 4.4 and pointed in obs. 2, most robotic projects currently use ROS, which uses a component-based approach. Therefore, robotic functionalities are often packaged as components that must be integrated into a robotic system. Integration issues are a consequence of such systems. From 13 provided responses with solutions, most respondents (seven) suggest to adhere to a clearly defined architecture and to model and document existing components and interfaces.

Simulation-to-Real-World Transition. This issue is considered a major challenge by 58% of the survey respondents. Interviewees and respondents claim that current simulation tools are not sufficient to properly validate robotic applications due to issues such as badly-simulated phenomena like physics and sensor feed. Among 13 collected solutions, five respondents opt for improving the simulation system based on machine learning techniques. Another five respondents suggest transitioning to real-world experimentation as soon as feasible. S14: “The simulation to real-world gap is real, don’t trust any performance measures of anything in simulation, especially if the system is complex. Transition to real world components as soon as it’s monetarily feasible.” Regarding this topic, we found:

Current simulation solutions are not able to emulate real-world phenomena in a sufficiently realistic manner. Therefore, practitioners distrust them for validating reliable robotic systems and instead they are relegated to validate basic robotic functionalities.

— Observation 14 (Real World to Simulation Transition) —

Safety Certification. More than half of the respondents (58%) considered safety certification a major challenge. As robots interact with the environment, and sometimes with humans, safety becomes a crucial issue. I17: “It’s hard to guarantee that a system that relies on learning from a data set, actually generalizes well [...] and learns the right things to be safe, once you deploy it on a completely different environment”. The 10 solutions obtained for this challenge are rather heterogeneous. For instance, one respondent employs a specialized team for such issues. Another respondent partners with a company with expertise in the domain. A third respondent trains their staff on related standards, such as ISO 13849 and IEC 61508/62061.

Lack of Standardized Solutions. 57% of the respondents agreed with this topic being a challenge. According to four of them, the lack of standardized solutions in robotics is directly related with its early state of maturity. S13: “The lack of standards, common datasets and problems, and general lack of reuse of software/infrastructure/APIs is making steady progress difficult.” We collect 14 answers with solutions for this topic. The one specified the most (named five times) is to develop internal standardized solutions as company guidelines. Two respondents opt to cherry-pick from solutions from other domains (e.g., A-SPICE [54] from the automotive domain).

Software Reuse. Among the respondents, 52% agreed or strongly agreed with software reuse being a challenge. I8: “In the future, implementing everything from scratch would get harder and harder. So the challenges that I see are basically how to reuse existing solutions and how to integrate them into your own development.” We discussed types, artifacts, and roadblocks to reuse in Sec. 4.6 (obs. 5, 6). We obtained 17 responses with solutions to this challenge. Six survey respondents encourage modularization and use of libraries and toolboxes. Other proposed solutions include adherence to standards, use of shared repositories, and model-driven development.

The lack of standard methods, frameworks, libraries, and reference models complicates the development of robotics systems. This, together with a large number of existing poorly-documented resources impede interoperability and easy reuse of software, which hinders the development pace of robotics SE.

— Observation 15 (Lack of Documentation and Established Conventions) —

Specifying Missions. Almost half of the respondents (48%) consider mission specification a challenge, being a key issue to make it easy enough for end-users. Also, mission specification mechanisms should enable users to express “exceptional” behaviors to specify how robots ought to cope with the variability of conditions of application scenarios in real environments in which robots are required to operate [26]. This is especially problematic in scenarios involving humans (obs 11). We collected 10 responses with solutions, divided into two groups: mission specification with and without end-user involvement. If the end-user is directly involved, methods including DSLs, GUIs, and frameworks are proposed. If not, developers need to “hard code” the mission into the system, typically based on iterative meetings with the end-users. End-users need to be involved

in this iterative mission specification in the case of multi-purpose robots that might be used for a variety of missions and/or scenarios that cannot be or it is impractical to specify upfront.

Applying AI Techniques. This topic was considered a challenge by 34% of the respondents. Out of 18 interviewees, seven apply AI techniques to their projects, five of them to perform obstacle and object detection. We collected three solutions to this challenge, including applying and extending current research to solve relevant problems, and to provide additional education.

Others. Of “Other” responses, 90% answered “Don’t know.” The remaining 10% provided eight responses, including: “deterministic execution, timestamping, clock synchronization,” “self-x capabilities,” and “lacking extensibility or clear APIs leading to forking.”

7 IMPACT AND PERSPECTIVES

Our collected data and the answers to our three research questions give rise to the following three research themes deserving further attention. We discuss them along with observations, formulate hypotheses trying to explain them, and propose actions.

7.1 Robustness and System Validation

Observations. Even though testing practices are among the most widespread quality assurance techniques in robotics (obs. 3), and thorough real-world testing and experimentation processes are mentioned as the most common activities performed by our respondents (obs. 1), our results point out that robustness and validation are still perceived as the top challenges of the domain (obs. 13). The respondents emphasize the importance of real-world testing especially in the light of the inaccuracy of simulations (obs. 14, discussed below). Respondents also highlight the difficulty and the importance of programming robots to handle failures (obs. 13), i.e., act robustly.

Hypotheses. We hypothesize that the importance of robustness and real-world testing relates to the fact that service robots are required to operate without human intervention in a large variety of real-world contexts, which are characterized by an unwieldy diversity of sources of uncertainty. Systematic testing processes can support the management of possible failures and unexpected environmental events over prolonged timeframes.

Proposed actions. As real-world testing is expensive, practitioners would largely benefit from techniques that provide realistic test data. Most contemporary methods in the testing domain generate test data by relying on code-centric information like coverage criteria and input spaces. In contrast, there is a potential for novel testing techniques that consider historical real-world data and use them for systematically producing new critical scenarios, which we propose as an action for researchers. Collections of real-world, industry-based critical scenarios are required to recreate testing cases and therefore improve the robustness of robotic systems. These critical scenarios should be collected by industrial practitioners. Moreover, the uncertainty to which robots are exposed asks for novel mechanisms of instilling robustness to robots, e.g. domain-specific languages especially tailored to code environment-based exceptions so to cope with uncertainty. For the development of such mechanisms, we encourage collaboration among researchers and practitioners, being the former those who take the lead in the idea conception and prototype development and the latter those who apply their knowledge to sharpen the prototype into a refined tool.

7.2 Simulation

Observations. Since real-world experimentation is time demanding and costly, it is tempting to assume that simulation is the primary means of validating robotic functionalities. However, we found that practitioners rely on experimentation as much and sometimes even more than on simulation (obs. 1 and 4). This is due to the difficulty of modeling real-world phenomena and events typical of environments where robots operate, especially those that involve humans (obs. 11, 12, and 14). Therefore, improving the capabilities of robotics simulations is a key challenge.

Hypotheses. The insufficient ability of simulations to validate robotic systems trustworthily can be explained with the hypothesis that existing simulations are decoupled from the real world. While simulations are typically inspired by real-world phenomena, there is usually little effort to update them with real-world data to make them consistent with real-world observations. Doing so is a big challenge due to the complexity of the real world (see discussion above).

Proposed actions. Efforts to improve simulation may benefit from *Multi-Paradigm Modeling* (MPM, [53]), which suggests the use of appropriate formalisms and abstraction levels to model all relevant aspects of the system (e.g., the mission to perform, the robot embodiment) and the environment (e.g., human behavior, physics). The application of appropriate abstraction levels permits to hide the underlying complexity of those aspects, which are modeled with the aim of being usable by domain experts. Using MPM, those models from different domains are then automatically mapped to a common domain, which can be easily managed by developers [52]. Furthermore, MPM has been proposed as an effective manner of tackling the challenges of interdisciplinary domains design as is the example of cyber-physical systems [52]. We foresee a huge potential in the application of MPM techniques in robotics to allow domain experts to easily specify their concerns and model specific aspects that can then be integrated into a common domain. To bridge the gap between the physical and the virtual world in all life cycle phases, MPM could benefit from the use of *digital twin* technology [6], which allows keeping the models synchronized through the development life-cycle based on real-world data. We recommend close collaborations of researchers and practitioners to develop improved “live” MPM frameworks, where the former would act as the lead on the research and development and the latter can help by providing realistic models from their experimentation data.

7.3 Conventions, Frameworks, and Reference Models

Observations. The development of robotic software mostly consists of integrating reusable components that implement recurrent functionalities. While de facto conventions for hardware interfaces in some technical scopes and inter-component communication exist (e.g., ROS messages, see obs. 2), reusable solutions for recurrent problems in application development are still missing (obs. 15). In practice, integration is typically performed in an ad hoc way (obs. 7) while being complicated by the wide range of disciplines practitioners working to develop a full-working robotic system have (obs. 9). A related issue is the lack of documentation of implemented solutions, one of the main causes for the lack of reuse of existing resources (obs. 6). There have been attempts at providing universal solutions

for robotics, such as the OMG Robotics Domain Task Force [37], which aims at fostering the integration of robotic systems. However, to our knowledge, the latter has not produced relevant results so far.

Hypotheses. We hypothesize that the lack of success of such initiatives comes from the inherent essential complexity of the domain. Existing conventions on hardware interfaces are not global for the robotics domain, but specific to certain technical scopes, which does not adequately address the continuous supply with new types of hardware. Also, the robotics domain as a whole is divided into a variety of sub-domains, including vertical ones (e.g., drivers, planning, navigation) and horizontal ones (e.g., defense, healthcare, logistics).

Proposed actions. A common recommendation [12, 27] is to exploit and reuse reference architectures, instead of single components to ease the instantiation, implementation, and reuse of *system abilities* [46]. However, due to the domain's complexity, developing a standardized reference architecture for robotics might be infeasible. Instead, we propose to identify a reference architecture for each horizontal domain represented as “end-user market domains” in the Multi-Annual Robotics Roadmap [46]. Yet, these domains are subject to huge variability and it would be required to decompose them into specific, narrow sub-domains that can be used to define reference architectures. When focusing on a single sub-domain, the reference architecture should be profitably decomposed in various views and viewpoints, as recommended by the ISO/IEC/IEEE 42010 standard [30]. The standard recommends taming the complexity of modern systems through the use of various views, each conforming to a viewpoint, which is identified to frame specific concerns of a precise set of stakeholders. Then, the various viewpoints are contained into an architecture framework, which defines conventions, principles, and practices for the description of architectures established within a specific domain of application and/or community of stakeholders. Two studies of Pelliccione et al. are examples of an architecture framework in the automotive field [39], and of a specific viewpoint on connected vehicles [38]—the latter contains also a functional reference architecture for the specifically connected vehicle viewpoint.

Researchers in the software architecture community might bring expertise acquired in other fields to help to establish conventions, principles and practices, and developing architecture frameworks and reference architectures. To do so, researchers will need to collaborate with robotics practitioners. This may bring an important innovation in the field as it happened in the past with the arrival of middleware frameworks like ROS and OROCOS, which enormously simplified the reuse of software resources by (mainly) establishing communication architectures and interfacing conventions. This caused a shift in the community mindset, which understood the advantage of those frameworks to solve issues related to system integration—e.g., concurrency, synchronization, real-time communication. We suggest making similar efforts towards the standardization and establishment of conventions, principles, and practices.

8 THREATS TO VALIDITY

We discuss the threats to the validity using the standard categorization by Wohlin et al. [55].

External Validity. To ensure a high degree of external validity, we used a diverse recruitment strategy for our interviews and survey, described in Sec. 3. Our interviewees and respondents work with robotic systems, covering a wide range of domains, contexts,

and backgrounds. The practices, characteristics, and challenges we describe are applicable to service robots in similar domains.

Construct Validity. The main threat are terminology issues: while we benefit from our own experience in the domain, a lack of a shared understanding with our respondents may threaten our conclusions. As one example, we point out the case of requirements engineering (RE): only 49% of all survey respondents specify that RE is performed very often or always in their organizations. We noticed that several interviewees had a somewhat narrow understanding of RE, focusing on specialized notation and tools, and that a RE process in a wider sense was universally applied. Apart from this case, by collecting and triangulating qualitative and quantitative data, we were able to observe that the general trends were in line in both data sources.

Internal Validity. We relied on voluntary participants from industry, most of whom we did not know personally. Therefore, our sample may be affected by selection bias: it may be skewed towards participants who are willing to participate in volunteering activities, which may be correlated with some relevant phenomena such as the willingness to contribute to open-source projects.

Conclusion Validity. Coding as we performed it for our qualitative analysis is subject to interpretation. We mitigated bias by cross-checking the obtained codes and refining the coding in a collaborative manner. We triangulated the coding with our qualitative and quantitative data to enhance our study's validity.

9 CONCLUSION

We presented an empirical assessment of the current state of robotics SE, comprising 18 interviews with industrial practitioners from 15 different companies and a survey with 156 respondents from 26 countries. We discussed (1) common practices applied by practitioners for developing robotic applications, (2) the main characteristics of robotics SE, and (3) challenges commonly faced by practitioners and applied solutions. Robotic systems are required to provide higher levels of autonomy in terms of decision making and planning because they must interact with (and react to) the environment and humans. Software development for such systems renders important challenges and calls for systematic SE approaches to build them. During our study, we found an overall lack of SE best practices in robotics development (obs. 10) and a substantial awareness of this issue. In our understanding, robotics would enormously benefit from the application of such practices to solve some of the challenges we elicited—e.g., software architectures to cope with software and hardware interoperability, modeling to promote modularity and reusability. We hope our work contributes to raising awareness of the importance of further investigation of this topic.

Based on our discussion, possible directions for future work are: (1) Investigating novel testing techniques that consider real-world data; (2) the development of “live” multi-paradigm modeling frameworks that use the digital twin technology; and (3) the conception of reference architectures for specific end-user market domains.

ACKNOWLEDGMENTS

Research supported by the EU H2020 Prog. under GA No. 731869 (Co4Robots) and the Centre of EXcellence on Connected, Geo-Localized, and Cybersecure Vehicle (EX-Emerge), funded by Italian Government under CIPE resolution n. 70/2017 (Aug. 7, 2017). We are thankful to the participants for their collaboration.

REFERENCES

- [1] Aakash Ahmad and Muhammad Ali Babar. 2016. Software Architectures for Robotic Systems: A Systematic Mapping Study. *Journal of Systems and Software* 122 (2016), 16–39.
- [2] Harald Altinger. 2017. *State of the Art Software Development in the Automotive Industry and Analysis Upon Applicability of Software Fault Prediction*. Cuvillier Verlag.
- [3] Scientific American. 2016. Hitomi Spacecraft. <https://www.scientificamerican.com/article/software-error-doomed-japanese-hitomi-spacecraft/>.
- [4] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. 2019. Software Engineering for Machine Learning: A Case Study. In *ICSE-SEIP*. IEEE Press, 291–300.
- [5] Saddek Bensalem, Lavindra de Silva, Félix Ingrand, and Rongjie Yany. 2011. A Verifiable and Correct-by-Construction Controller for Robot Functional Levels. *Journal of Software Engineering for Robotics* 2 (2011), 9–19.
- [6] Stefan Boschert and Roland Rosen. 2016. Digital twin—the simulation aspect. In *Mechatronic futures*. Springer, 59–74.
- [7] Darko Bozhinoski, Davide Di Ruscio, Ivano Malavolta, Patrizio Pelliccione, and Ivica Crnkovic. 2019. Safety for mobile robotic systems: A systematic mapping study from a software engineering perspective. *Journal of Systems and Software* 151 (2019), 150–179.
- [8] Davide Brugali. 2010. From the Editor-in-Chief: A New Research Community, a New Journal. *Journal of Software Engineering for Robotics* 1 (2010), 1–2.
- [9] Davide Brugali. 2015. Model-Driven Engineering in Robotics. *IEEE Robot. Automat. Mag.* 22, 3 (2015), 155–166.
- [10] Davide Brugali, Arvin Agah, Bruce MacDonald, Issa A. D. Nesnas, and William D. Smart. 2007. *Trends in Robot Software Domain Engineering*. Springer Berlin Heidelberg, Berlin, Heidelberg, 3–8.
- [11] Davide Brugali, Alex Brooks, Anthony Cowley, Carle Côté, Antonio C. Dominguez-Brito, Dominic Létourneau, François Michaud, and Christian Schlegel. 2007. *Trends in Component-Based Robotics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 135–142.
- [12] Davide Brugali, Luca Gherardi, A Biziak, Andrea Luzzana, and Alexey Zakharov. 2012. A reuse-oriented development process for component-based robotic systems. In *SIMPAR*. Springer, 361–374.
- [13] Davide Brugali and Erwin Prassler. 2009. Software engineering for robotics [From the Guest Editors]. *IEEE Robotics & Automation Magazine* 16, 1 (2009), 9–15.
- [14] Davide Brugali and Monica Reggiani. 2005. Software stability in the robotics domain: issues and challenges. In *IRL*. IEEE Press, 585–591.
- [15] Herman Bruyninckx, Markus Klotzbücher, Nico Hochgeschwender, Gerhard Kraetzschmar, Luca Gherardi, and Davide Brugali. 2013. The BRICS Component Model: A Model-based Development Paradigm for Complex Robotics Software Systems. In *SAC*. ACM, 1758–1764.
- [16] Herman Bruyninckx, Peter Soetens, and Bob Koninckx. 2003. The real-time motion control core of the Orocos project. In *ICRA*, Vol. 2. IEEE, 2766–2771.
- [17] Tomas Bures, Alessia Knauss, Pankesh Patel, Awais Rashid, Ivan Ruchkin, Roykrong Sukkerd, Christos Tsigkanos, Danny Weyns, Bradley Schmer, Eduardo Tovar, Eric Boden, Thomas Gabor, Ilias Gerostathopoulos, Pragma Gupta, and Eunsuk Kang. 2017. Software Engineering for Smart Cyber-Physical Systems: Challenges and Promising Solutions. *ACM SIGSOFT Software Engineering Notes* 42 (2017), 19–24.
- [18] Tomas Bures, Danny Weyns, Christian Berger, Stefan Biffl, Marian Daun, Thomas Gabor, David Garlan, Ilias Gerostathopoulos, Christine Julien, Filip Krikava, et al. 2015. Software engineering for smart cyber-physical systems—towards a research agenda: report on the First International Workshop on Software Engineering for Smart CPS. *ACM SIGSOFT Software Engineering Notes* 40, 6 (2015), 28–32.
- [19] The Construct. 2019. Top 10 ROS Based Robotics Companies in 2019. <http://www.theconstructsim.com/top-10-robotics-companies-2019-ros-based>.
- [20] Juliet M Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative sociology* 13, 1 (1990), 3–21.
- [21] CORDIS. 2013. Best Practice in Robotics (BRICS). <https://cordis.europa.eu/project/id/231940>.
- [22] CORDIS. 2020. Co4Robots. <https://cordis.europa.eu/project/id/731869>.
- [23] Ayssam Elkady and Tarek Sobh. 2012. Robotics Middleware: A Comprehensive Literature Survey and Attribute-Based Bibliography. *Journal of Robotics* (2012).
- [24] Pablo Estefó, Jocelyn Simmonds, Romain Robbes, and Johan Fabry. 2019. The Robot Operating System: Package reuse and community dynamics. *Journal of Systems and Software* 151 (2019), 226–242.
- [25] Sergio García, Daniel Strüber, Davide Brugali, Thorsten Berger, and Patrizio Pelliccione. 2020. Replication Package. Replication package website: <https://sites.google.com/view/empirical-study-robotics-se/home> Archived at: https://figshare.com/articles/conference_contribution/Robotics_Software_Engineering_A_Perspective_from_the_Service_Robotics_Domain/12376931 With DOI: 10.6084/m9.figshare.12376931.
- [26] Sergio García, Daniel Strüber, Davide Brugali, Alessandro Di Fava, Philipp Schillinger, Patrizio Pelliccione, and Thorsten Berger. 2019. Variability Modeling of Service Robots: Experiences and Challenges. In *VaMoS*. ACM, 8.
- [27] Luca Gherardi and Davide Brugali. 2014. Modeling and reusing robotic software architectures: the hyperflex toolchain. In *ICRA*. IEEE, 6414–6420.
- [28] Félix Ingrand and Malik Ghallab. 2017. Deliberation for autonomous robots: A survey. *Artificial Intelligence* 247 (2017), 10–44.
- [29] ISO. 2012. ISO 8373:2012. Robots and robotic devices - Vocabulary. <https://www.iso.org/standard/55890.html>.
- [30] ISO/IEC/IEEE. 2011. ISO/IEC/IEEE 42010:2011. Systems and software engineering - Architecture description. <https://www.iso.org/standard/50508.html>.
- [31] Nathan Koenig and Andrew Howard. 2004. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IRIOS*, Vol. 3. IEEE, 2149–2154.
- [32] David Kortenkamp, Reid Simmons, and Davide Brugali. 2016. Robotic systems architectures and programming. In *Springer Handbook of Robotics*. Springer, 283–306.
- [33] Markets and Markets. 2017. Service Robotics Market – Global Forecast to 2022. <https://www.marketsandmarkets.com/Market-Reports/service-robotics-market-681.html>.
- [34] Yuya Maruyama, Shinpei Kato, and Takuya Azumi. 2016. Exploring the performance of ROS2. In *EMSOFT*. 1–10.
- [35] Giorgio Metta, Paul Fitzpatrick, and Lorenzo Natale. 2006. YARP: yet another robot platform. *International Journal of Advanced Robotic Systems* 3, 1 (2006), 8.
- [36] Arne Nordmann, Nico Hochgeschwender, Dennis Wigand, and Sebastian Wrede. 2016. A Survey on Domain-Specific Modeling and Languages in Robotics. *Journal of Software Engineering for Robotics* 7 (2016), 75–99.
- [37] OMG. 2019. Robotics Domain Task Force. <https://www.omg.org/robotics/>.
- [38] Patrizio Pelliccione, Eric Knauss, Magnus Ågren, Rogardt Heldal, Carl Bergenhem, Alexey Vinel, and Oliver Brunnegård. 2020. Beyond connected cars: A systems of systems perspective. *Science of Computer Programming* 191 (2020), 102414.
- [39] Patrizio Pelliccione, Eric Knauss, Rogardt Heldal, Magnus Ågren, Piergiuseppe Mallozzi, Anders Alming, and Daniel Borgentun. 2017. Automotive architecture framework: The experience of Volvo cars. *Journal of systems architecture* 77 (2017), 83–100.
- [40] Kari Pulli, Anatoly Baksheev, Kirill Korniyakov, and Victor Eruhimov. 2012. Real-time computer vision with OpenCV. *Commun. ACM* 55, 6 (2012), 61–69.
- [41] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. 2009. ROS: an open-source Robot Operating System. *ICRA workshop on open source software* 3, 3.2 (2009), 5.
- [42] Christian Schlegel. 2007. Communication patterns as key towards component interoperability. In *Software Engineering for Experimental Robotics*. Springer, 183–210.
- [43] Christian Schlegel, Thomas Haßler, Alex Lotz, and Andreas Steck. 2009. Robotic software systems: From code-driven to model-driven designs. In *ICRA*. IEEE, 1–8.
- [44] Douglas C. Schmidt. 2006. Guest Editor’s Introduction: Model-Driven Engineering. *Computer* 39, 2 (2006).
- [45] Forrest Shull, Janice Singer, and Dag IK Sjøberg. 2007. *Guide to advanced empirical software engineering*. Springer.
- [46] SPARC. 2016. Robotics 2020 Multi-Annual Roadmap. <https://eu-robotics.net/sparc/upload/about/files/H2020-Robotics-Multi-Annual-Roadmap-ICT-2016.pdf>.
- [47] Andreas Steck, Alex Lotz, and Christian Schlegel. 2011. Model-driven Engineering and Run-time Model-usage in Service Robotics. In *GPCE*. 73–82.
- [48] Anselm Strauss and Juliet Corbin. 1990. Open Coding. *Basics of Qualitative Research: Grounded Theory Procedures and Techniques* 2 (1990), 101–121.
- [49] SVR. 2019. Women in robotics on International Women’s Day 2019. <https://svrobo.org/women-in-robotics-on-international-womens-day-2019>.
- [50] Walter F Tichy. 2000. Hints for reviewing empirical work in software engineering. *Empirical Software Engineering* 5, 4 (2000), 309–312.
- [51] Hans Utz, Stefan Sablatnig, Stefan Enderle, and Gerhard Kraetzschmar. 2002. Miro-middleware for mobile robot applications. *IEEE Transactions on Robotics and Automation* 18, 4 (2002), 493–497.
- [52] Yentl Van Tendeloo and Hans Vangheluwe. 2017. The Modelverse: a tool for multi-paradigm modelling and simulation. In *WSC*. IEEE, 944–955.
- [53] Hans Vangheluwe, Juan De Lara, and Pieter J Mosterman. 2002. An introduction to multi-paradigm modelling and simulation. In *AIS*. 9–20.
- [54] VDA QMC Working Group 13 / Automotive SIG. 2015. Automotive SPICE Process Assessment / Reference Model. Issue 3.0.
- [55] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in software engineering*. Springer Science & Business Media.