Robot Cybersecurity, a review

Abstract

Robots are often shipped insecure and in some cases fully unprotected. The rationale behind is threefold: first, defensive security mechanisms for robots are still on their early stages, not covering the complete threat landscape. Second, the inherent complexity of robotic systems makes its protection costly, both technically and economically. Third, vendors do not generally take responsibility in a timely manner, extending the zero-days exposure window (time until mitigation of a zero-day) to several years on average. Worse, several manufacturers keep forwarding the problem to the end-users of these machines or discarding it.

In this article we review the status of robot cybersecurity considering three sources of data: 1) recent literature, 2) questionnaires performed in top robotics forums and 3) recent research results in robot cybersecurity. Building upon a decade of experiences in robotics, this article reviews the current status of cybersecurity in robotics and argues about the current challenges to secure robotic systems. Ultimately, based on the empirical results collected over a period of three years performing security assessments in robots, the present text advocates for a complementary offensive approach methodology to protect robots in a feasible and timely manner.

Notes for Practice

- This article reviews the status of robot cybersecurity considering research and empirical data.

- While systematically reviewing sources, the article provides various observations that attempt answering two questions posed initial: what's the status of cybersecurity in robotics? and, how can we best improve cyber-resillience in robotics?

- Following a red teaming methodology, the article considers the vulnerabilities affecting robots from three different robot manufacturers and provides an intuition on how to use these flaws to generate indicators to evaluate the security readiness of a manufacturer

- Based on the theoretical and empirical results studied, the text concludes that robot cybersecurity is still on its early stages and neither practitioners nor manufacturers or defensive security solutions meet the current industry demands to remain secure. Given this scenario the text provides a final observation considering red teaming methodologies to approach robot cybersecurity in a time-sensitive manner.

1. Introduction

For the last fifty years, we have been witnessing the dawn of the robotics industry, but robots are not being created with security as a concern, often an indicator of a technology that still needs to mature. Security in robotics is often mistaken with safety. From industrial to consumer robots, going through professional ones,

ISSN 2753-9997 (online). The International Journal of Cyber Forensics and Advanced Threat Investigations under a Creative Commons License, Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0)

1

most of these machines are not resilient to security attacks. Manufacturers' concerns, as well as existing standards, focus mainly on safety. Security is not being considered as a primary relevant matter.

The integration between these two areas from a risk assessment perspective was first studied by Stoneburner (2006) and later discussed by Alzola-Kirschgens et al. (2018) which resulted in a unified security and safety risk framework. Commonly, robotics safety is understood as developing protective mechanisms against accidents or malfunctions, whilst security is aimed to protect systems against risks posed by malicious actors (Swinscow-Hall, 2017). A slightly alternative view is the one that considers safety as protecting the environment from a given robot, whereas security is about protecting the robot from a given environment. In this article we adopt the latter and refer the reader to the appendix A for a more detailed literature review that introduces the differences and correlation between safety, security and quality in robotics.

Security is not a product, but a process that needs to be continuously assessed in a periodic manner, as systems evolve and new cyber-threats are discovered. This becomes specially relevant with the increasing complexity of such systems as indicated by Bozic and Wotawa (2017). Current robotic systems are of high complexity, a condition that in most cases leads to wide attack surfaces and a variety of potential attack vectors which makes difficult the use of traditional approaches. Altogether, this leads to the following research questions: what's the status of cybersecurity in robotics? and, how can we best improve cyber-resilience in robotics?

The present article tackles these questions by performing a systematic review. The approach followed is three-fold: first, we review literature in the robot cybersecurity space depicting the current landscape. Second, we study and review the results obtained while surveying different robotics groups and communities in search for an answer of the current state of robot cybersecurity. Third and lastly, we review the results obtained during three years of proactive security research in robotics, discussing vulnerabilities and the offensive exercises conducted. The article finalizes by sharing some thoughts and conclusions about how to secure robots, how to understand better the attack vectors (they are subject to) and how to minimize their attack surfaces.

Robotic systems and robots

Both literature and practice are often vague when using the terms *robot*/s and/or *robotic system*/s. Sometimes these terms are used to refer to one of the robot components (e.g. the robot is the robot arm mechanics while its human-machine interface (HMI) is the *teach pendant*). Some other times these terms are used to refer to the complete robot, including all its components, regardless of whether they are distributed or assembled into the same hull. Throughout this article the latter is adopted and unless stated otherwise, the terms *robot*/s and/or *robotic system*/s will be used interchangeably to refer to the complete robotic system, including all its components.
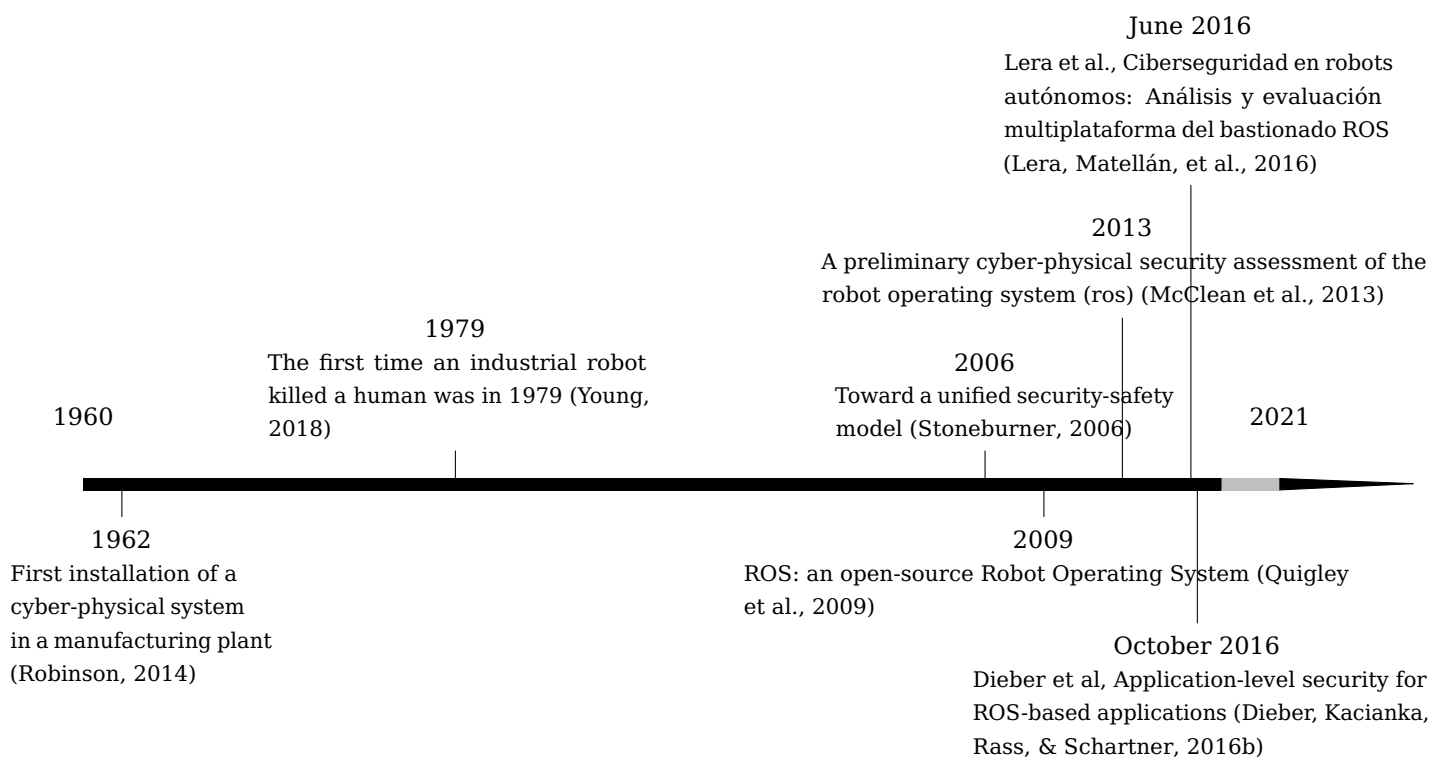
The remainder of the article is organized as follows: section 2 provides a literature review of the robot cybersecurity field highlighting some of the current biographical cornerstones. Section 3 elaborates on the results of the questionares and surveys performed in top robotics conferences and forums. Section 4 reviews the vulnerability landscape in robotics using empirical data collected while assessing the security of multiple robots over the past three years. Finally, section 5 presents some conclusions and thoughts on how to better secure robotic systems.

2. A literature review of cybersecurity in robotics

Arguably, the first installation of a cyber-physical system in a manufacturing plant was back in 1962 (Robinson, 2014). The first human death caused by a robotic system is traced back to 1979 (Young, 2018) and the causes were safety-related according to the reports. From this point on, a series of actions involving agencies and corporations triggered to protect humans and environments from this machines, leading into safety standards.
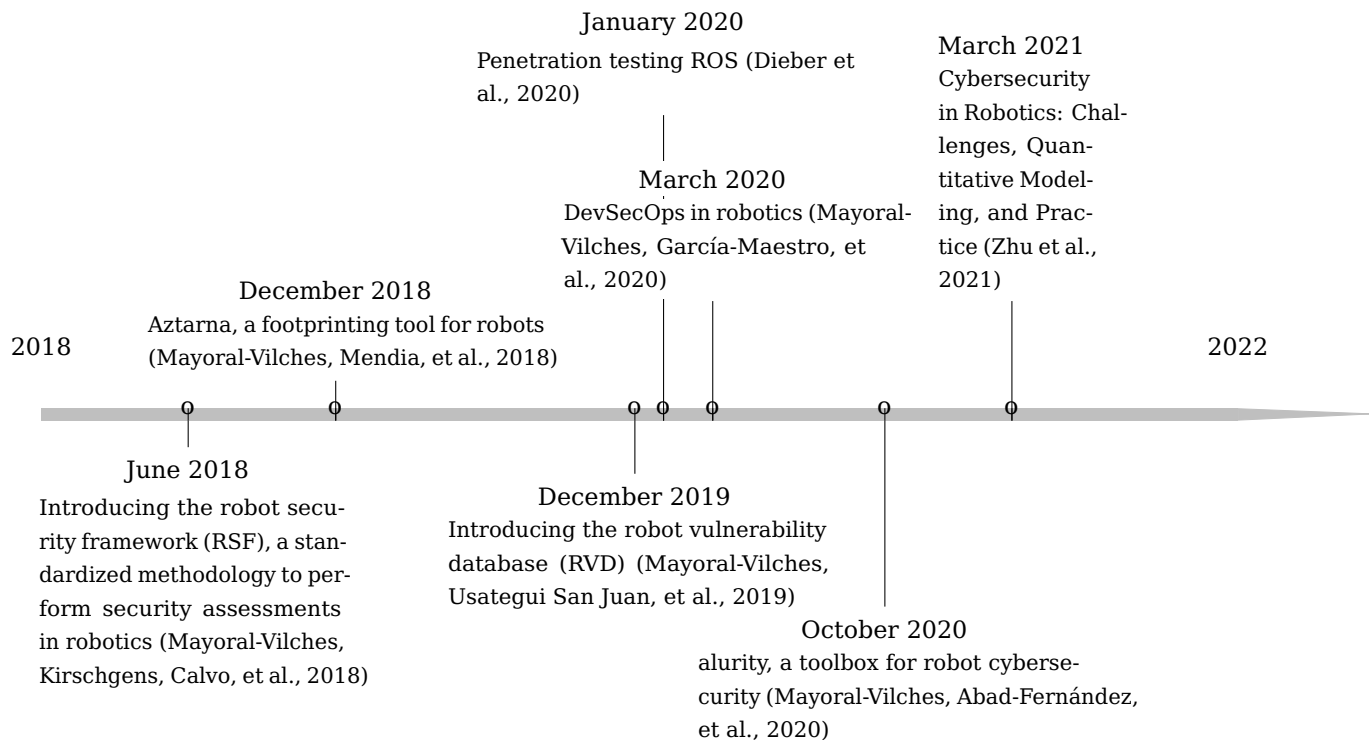
Security however hasn't started being addressed in robotics until recently. Following after McClean et

al. (2013) early assessment, in one of the first published articles on the topic Lera, Matellán, et al. (2016) already warns about the security dangers of the Robot Operating System (ROS) (Quigley et al., 2009). Following from this publication, the same group in Spain authored a series of articles touching into robot cybersecurity (Lera, Balsa, et al., 2016; Lera, Llamas, Guerrero, & Olivera, 2017; Guerrero-Higueras, DeCastro-García, Rodríguez-Lera, & Matellán, 2017; Balsa-Comerón, Guerrero-Higueras, Rodríguez-Lera, Fernández-Llamas, & Matellán-Olivera, 2017; Rodríguez-Lera, Matellán-Olivera, Balsa-Comerón, Guerrero-Higueras, & Fernández-Llamas, 2018). Around the same time period, Dieber et al. (2016a) led a series of publications that researched cybersecurity in robotics proposing defensive blueprints for robots built around ROS (Dieber, Breiling, et al., 2017; Dieber, Schlotzhauer, & Brandstötter, 2017; Breiling, Dieber, & Schartner, 2017; Taurer, Dieber, & Schartner, 2018; Dieber & Breiling, 2019). Their work introduced additions to the ROS APIs to support modern cryptography and security measures. Contemporary to Dieber et al. (2016a)'s work, White et al. (2016) also started delivering a series of articles (Caiazza, 2017; White, Christensen, Caiazza, & Cortesi, 2018; White, Caiazza, Christensen, & Cortesi, 2019; Caiazza, White, & Cortesi, 2019; White, Caiazza, Jiang, et al., 2019; White, Caiazza, Cortesi, Im Cho, & Christensen, 2019) proposing defensive mechanisms for ROS.

June 2016
Lera et al., Ciberseguridad en robots autónomos: Análisis y evaluación multiplataforma del bastionado ROS (Lera, Matellán, et al., 2016)

2013
A preliminary cyber-physical security assessment of the robot operating system (ros) (McClean et al., 2013)

1979
The first time an industrial robot killed a human was in 1979 (Young, 2018)

2006
Toward a unified security-safety model (Stoneburner, 2006)

1960

2021

1962
First installation of a cyber-physical system in a manufacturing plant (Robinson, 2014)

2009
ROS: an open-source Robot Operating System (Quigley et al., 2009)

October 2016
Dieber et al, Application-level security for ROS-based applications (Dieber, Kacianka, Rass, & Schartner, 2016b)

A bit more than a year after that, starting in 2018, it's possible to observe how more groups start showing interest for the field and contribute. Mayoral-Vilches, Kirschgens, Calvo, et al. (2018) initiated a series of security research efforts attempting to define offensive security blueprints and methodologies in robotics that led to various contributions (Mayoral-Vilches, Kirschgens, Gil-Uriarte, et al., 2018; Alzola-Kirschgens et al., 2018; Mayoral-Vilches, Mendia, et al., 2018; Mayoral-Vilches, Abad-Fernández, et al., 2020; Mayoral-Vilches, Pinzger, et al., 2020; Lacava et al., 2020; Mayoral-Vilches, García-Maestro, et al., 2020; Mayoral-Vilches, Carbajo, & Gil-Uriarte, 2020). Most notably, this group released publicly a framework for conducting security assessments in robotics (Mayoral-Vilches, Kirschgens, Calvo, et al., 2018), a vulnerability scoring mechanism for robots (Mayoral Vilches et al., 2018), a capture the flag (CTF) environment for robotics whereto learn how to train robot cybersecurity engineers (Mendia et al., 2018) or a robot-specific vulnerability database that third parties could use to track their threat landscape (Mayoral-Vilches, Usategui San Juan, et al., 2019),

among others. In 2021, Zhu et al. (2021) published a comprehensive introduction of this emerging topic for theoreticians and practitioners working in the field to foster a sub-community in robotics and allow more contributors to become part of the robot cybersecurity effort.

January 2020
Penetration testing ROS (Dieber et al., 2020)

March 2021
Cybersecurity in Robotics: Challenges, Quantitative Modeling, and Practice (Zhu et al., 2021)

March 2020
DevSecOps in robotics (Mayoral-Vilches, García-Maestro, et al., 2020)

December 2018
Aztarna, a footprinting tool for robots (Mayoral-Vilches, Mendia, et al., 2018)

2018

2022

June 2018
Introducing the robot security framework (RSF), a standardized methodology to perform security assessments in robotics (Mayoral-Vilches, Kirschgens, Calvo, et al., 2018)

December 2019
Introducing the robot vulnerability database (RVD) (Mayoral-Vilches, Usategui San Juan, et al., 2019)

October 2020
alurity, a toolbox for robot cybersecurity (Mayoral-Vilches, Abad-Fernández, et al., 2020)

A careful review of the prior art described in the last paragraphs leads to the following observation:

Observation 1 *Based on literature, robot cybersecurity is still a new field that deserves further attention, tools and educational material to train new engineers in security practices for robotics.*

3. Surveying the robotics community

During a period of two years (2019 - 2020) various security surveys were conductued in top robotics conferences and forums. The following subsections discuss each one of them while attempting to draw some observations:

3.1 Surveying the ROS community

Figure 1 presents a summarized result of the survey conducted in the ROS community during a period of several months. The survey received a total of 52 responses, which represented the small interest in security at the time. The largest groups of participants are depicted in Figure 1b. The most represented group comes from Universities (30%), followed by Software vendors (18%) and Robot manufacturers (14%)[1]. The majority of the respondents have at last 2 years of experience with ROS and half of them at least 5 (1c), most coming from Europe (1d). Figure 1e present data on security considerations. The data indicates that 73% of the participants think that they have not invested enough to protect their robots from cyber-threats. Coincidentally, the same

---

[1]Others comprises various subgroup, all with less representation than the ones mentioned

number of participants indicated that their organizations are open to invest however only 26% acknowledge to actually have invested. This data leads to the following observation:

Observation 2 *There's a gap between the expectations and the actual investment, which suggests that cybersecurity actions in robotics will grow in the future for the ROS community.*

When considering the mitigation strategies applied by respondents as depicted in Figure 1f, it's important to highlight that most efforts concentrate on perimeter actions (i.e. firewalls, segmentation and segregation) whereas robot-specific defensive solutions are only applied in a 36% of the cases. Similarly, network assessments and security audits are conducted only in one fourth of the cases (26%) which conflicts with the *de facto* security practices in other industries, wherein assessments are critical to evaluate the resilience of technology.

Observation 3 *The lack of robot-specific security measures (36%) and offensive assessments (26%) can be interpreted as an indicator of the maturity level of the technology when compared to other sectors (e.g. IT or OT) where these practices are common and specialized.*

3.2 Surveying the PX4 community

PX4 (Meier et al., 2015) is an open source flight control software for drones and other unmanned vehicles. Similar to ROS, its community represents another relevant group in robotics. A security survey was conducted in 2020 and the results are summarized in Figure 2. Though the PX4 community is significantly smaller than ROS's, the sample size obtained (11 respondents) was extremely small to draw major conclusions. Interestingly though, it was observed that the majority of the respondents have yet to see a security issue impacting the community (2d), only 27% had seen it.

Observation 4 *Both the PX4 (figure 2d) and the ROS (figure 1e) communities indicated that the majority is yet to witness a cyber-attack. In the ROS community only one out of ten respondents (9%) had seen it whereas in the PX4 group, approximately one out of four (27%).*

The majority of the respondents (81%, figure 2e) indicated to be willing to invest and more than 90% confirmed that the amount could be 100 USD or above (figure 2f). This aligns nicely with observation 2 and further hints that growth should be expected in this field.

3.3 Surveying the ROS-Industrial community

Also in 2020, a series of security-related surveys were launched as part of the European ROS-Industrial Conference, which happens every year in December. Data collected is presented in Figure 3. The majority of the respondents (93%) showed awareness about the threats their robots faced and admitted being aware of their exposure to attackers (figure 3b). Unsurprisingly, as a subset of the overall ROS community, the security mitigation actions in the ROS-I community also concentrate on the perimeter which lead to another observation:

Observation 5 *Figures 1f and 3d confirm that respectively for both ROS and ROS-I groups mitigations concentrate mostly on the perimeter.*

This fact becomes concerning in industrial environments wherein insider threats are as dangerous, and the disruption of ROS could lead to catastrophic consequences for the automation processes (Mayoral-Vilches, Pinzger, et al., 2020), impacting more than 5 robots in 44% of the cases according to respondents (figure 3f). The lack of security measures in ROS are particularly concerning since its distributed communication middleware could be easily used to spread malware across connected robots. Such concept was demonstrated by Mayoral-Vilches, San Juan, et al. (2019), which prototyped an instance of ransomware targeting industrial collaborative robots, leaving these machines and its data completely locked until the corresponding ransom is paid.
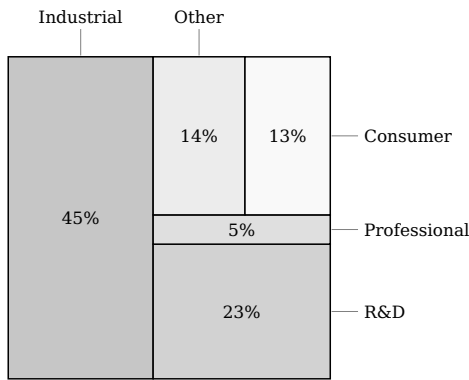
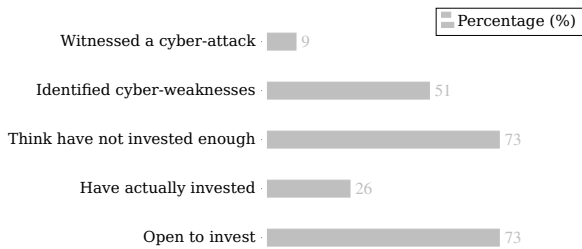**(a)** Distribution of respondents by sector of activity. *Sample size 52 respondents.*

**(b)** Distribution within the robotics value chain. *Sample size 52 respondents.*
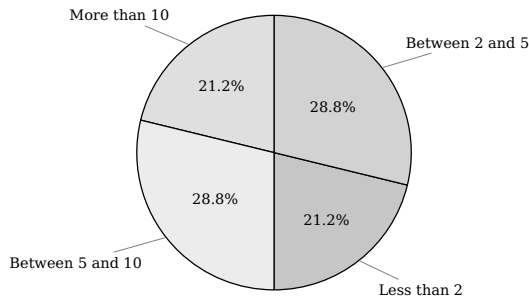
(c) Years of experience with ROS of each respondent. *Sample size 52 respondents.*

**(d)** Geographical distribution of respondents. *Sample size 52 respondents.*

**(e)** Percentage favouring each security consideration amongst respondents (in the context of robotics). *Sample size 52 respondents.*

(f) Percentages favouring each mitigation strategy amongst respondents (in the context of robotics). *Sample size 52 respondents.*

Figure 1. Surveying the ROS robotics community (2019): Security survey launched within the ROS Discourse community (announcement, announcement 2, preliminary results).

(a) Distribution of respondents to the question: "What's your cybersecurity concern in PX4?". No respondents indicated "Low". *Sample size 11 respondents.*


(b) Distribution of respondents to the question: "Does cyber security affect safety?". Other corresponds with user-provided answers. *Sample size 11 respondents.*


(c) Distribution of respondents to the question: "What's your GCS?". Various additional ground control stations (GCSs) were offered but not selected. *Sample size 11 respondents.*

(d) Distribution of respondents to the question: "Have you seen any security issues so far in PX4?". *Sample size 11 respondents.*


(e) Distribution of respondents to the question: "Would you be willing to invest in cybersecurity for your drone?". *Sample size 11 respondents.*


(f) Distribution of respondents to the question: "How much per year is security worth to you?". *Sample size 11 respondents.*

Figure 2. Surveying the PX4 robotics community (2020): Security survey launched within the PX4 Discourse community (announcement).
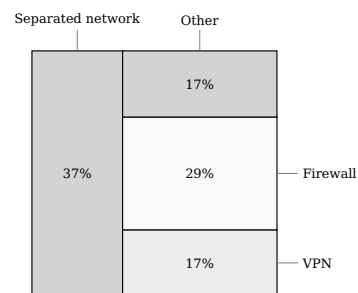
(a) Distribution of respondents to the question: "How important do you think is security for robotics and automation?". *Sample size 30 respondents.*


(b) Distribution of respondents to the question: "Do you think your robot can be hacked?". *Sample size 28 respondents.*


(c) Distribution of respondents to the question: "Have you taken measures to protect your robots?". *Sample size 25 respondents.*

(d) Distribution of respondents to the question: "What measures do you take to secure your robots?". *Sample size 23 respondents.*

(e) Distribution of respondents to the question: "Did you use fuzzing before?". *Sample size 27 respondents.*

(f) Distribution of respondents to the question: "How many robots are you controlling with ROS in your organization?". *Sample size 55 respondents.*

Figure 3. Surveying the ROS-Industrial robotics community (2020): Security surveys launched within the ROS-Industrial community during the digital ROS-I Europe Conference in December 2020.

3.4 Surveying the European robotics community at the European Robotics Forum (ERF) (2020)

As one of the leading geographies in robotics and cybersecurity, the opinion of european robotics experts was sampled during the annual European Robotics Forum (ERF). Figure 4 summarizes the most relevant data collected. The most interesting observation relates to the question *"Who is the actor to be responsible for cyber-incidents?"*:

Observation 6 *In Europe, the majority of the respondents (figure 4b) agree that the responsibility in case of damage as a result of a cyber-incident is to be assumed by the supply chain (86% indicated that it'd sit between System Integrators and robot vendors), with only a 14% pushing the responsibility to the end-user.*

4. Security research results in robotics

Figure 5 depicts summarized vulnerability research results for three vendors: ABB, Mobile Industrial Robots (MiR) and Universal Robots (UR). The data was collected and archived over a multi-year period. Figures 5a, 5b and 5c illustrate the "days until mitigation" for each vulnerability and according to the public data in the Robot Vulnerability Database (RVD) (Mayoral-Vilches, Usategui San Juan, et al., 2019). The flat line represented by a series of data points in figures 5b and 5c denotes that the vendor hasn't reacted yet to any of these flaws and they remain unmitigated (they are zero days). For ABB robots, the scattered plot in figure 5a denotes more security activity. The following observations are drawn from the data:

Observation 7 *Collaborative robot manufacturers MiR and UR have zero days with an age at least older than one year (figures 5b and 5c). These flaws continue growing older due to the inactivity from the manufacturers.*

Observation 8 *Vulnerability data affecting ABB robots (figure 5a) shows how according to historical data, vulnerabilities were patches as early as 14 days after its disclosure however the average mitigation time is above four years (1500 days).*

On top of these, Figures 5d to 5i enhance previous data with additional private sources of information and consider vulnerabilities that have yet to reach the public domain. It should be noted that the distribution of vulnerabilities signals the security awareness of the manufacturer. Coherently, figure 5g shows how for ABB robots, four out of five vulnerabilities considered have been publicly disclosed, triaged and scored. In contrast, for MiR and UR robots the oppositive is observed. Four out of five vulnerabilities have yet to be disclosed publicly.

Observation 9 *The ratio of publicly disclosed vulnerabilities versus the ones remaining private is an indicator when evaluating the security readiness of a robot manufacturer. The threat landscape of a given robot is correlated to this ratio in a direct manner.*
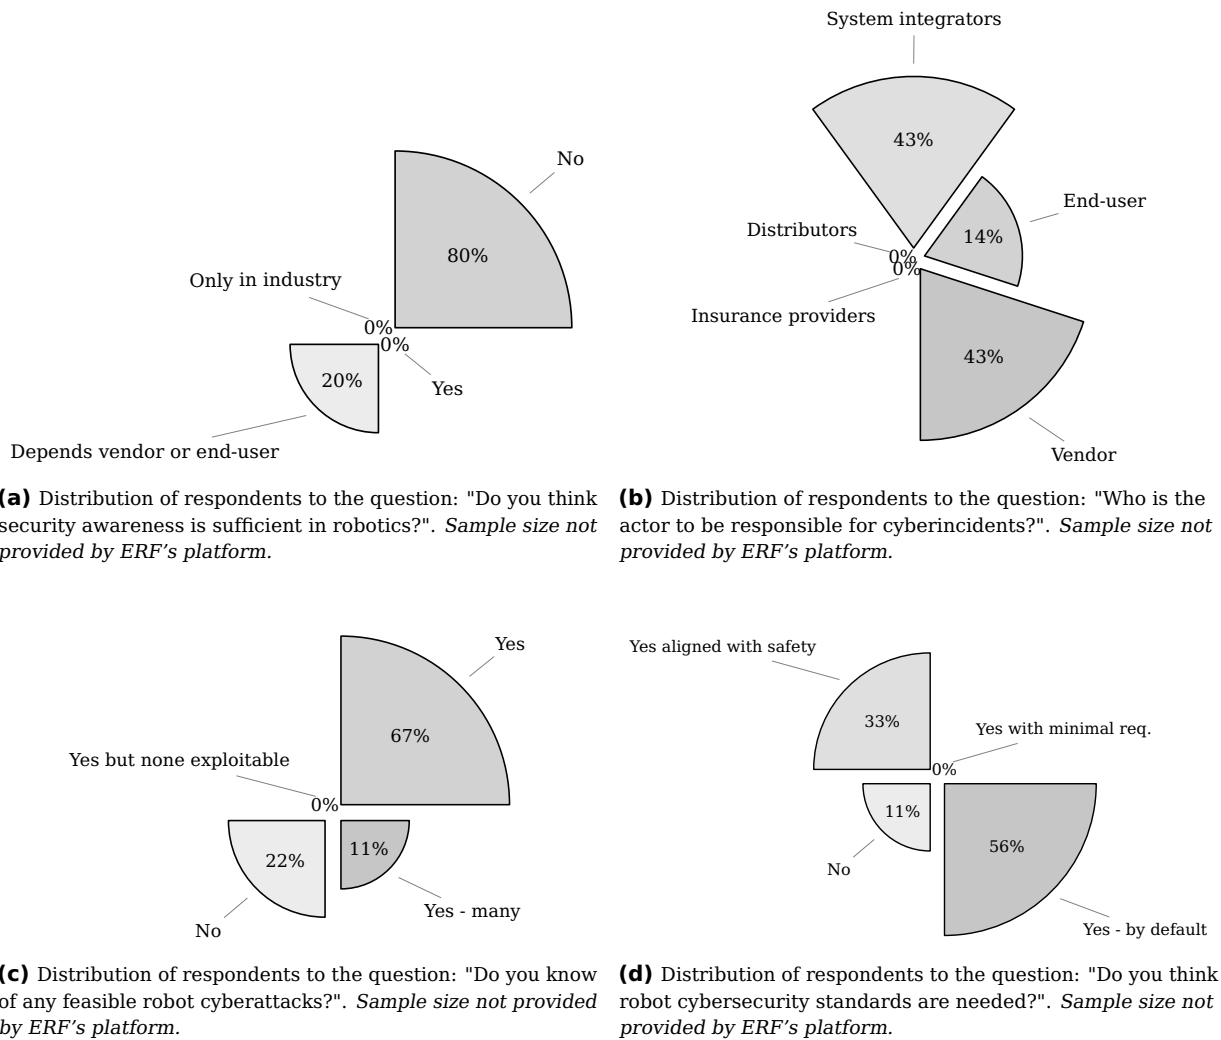
(a) Distribution of respondents to the question: "Do you think security awareness is sufficient in robotics?". *Sample size not provided by ERF's platform.*

(b) Distribution of respondents to the question: "Who is the actor to be responsible for cyberincidents?". *Sample size not provided by ERF's platform.*

(c) Distribution of respondents to the question: "Do you know of any feasible robot cyberattacks?". *Sample size not provided by ERF's platform.*

(d) Distribution of respondents to the question: "Do you think robot cybersecurity standards are needed?". *Sample size not provided by ERF's platform.*

Figure 4. Surveying the European robotics community (ERF 2020): Security surveys conducted during the robotics European gathering at the European Robotics Forum (ERF) 2020 in Málaga. The questionares were launched during the security sessions.

(a) Days until mitigation for each one of the vulnerabilities publicly registered for ABB robots. *Data collected from the RVD (Mayoral-Vilches, Usategui San Juan, et al., 2019).*

(b) Days until mitigation for each one of the vulnerabilities publicly registered for MiR robots. *Data collected from the RVD (Mayoral-Vilches, Usategui San Juan, et al., 2019).*

(c) Days until mitigation for each one of the vulnerabilities publicly registered for Universal Robots robots. *Data collected from the RVD (Mayoral-Vilches, Usategui San Juan, et al., 2019).*
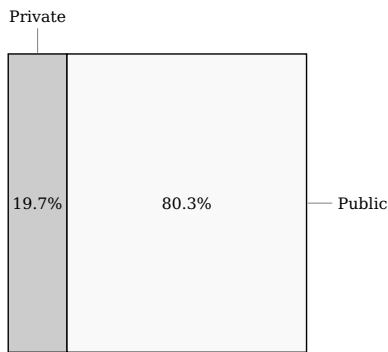
(d) Vulnerabilities affecting ABB robots registered in the RVD (Public), in other private databases (Private), as well as the overall amount (Total).

(e) Vulnerabilities affecting MiR robots registered in the RVD (Public), in other private databases (Private) as well as the overall amount (Total).

(f) Vulnerabilities affecting UR robots registered in the RVD (Public), in other private databases (Private) as well as the overall amount (Total).

(g) Distribution of the vulnerabilities affecting ABB robots and registered in the RVD (Public) or in other private databases (Private).

(h) Distribution of the vulnerabilities affecting MiR robots and registered in the RVD (Public) or in other private databases (Private).

(i) Distribution of the vulnerabilities affecting UR robots and registered in the RVD (Public) or in other private databases (Private).

Figure 5. Vulnerability data for various robots: the main source of public data is the Robot Vulnerability Database (RVD) Mayoral-Vilches, Usategui San Juan, et al. (2019). *Private data has been facilitated by Alias Robotics as an extension of RVD.*

5. Conclusions

This article reviews the current status of cybersecurity in robotics and argues about the current challenges to secure robotic systems. Three sources of data were considered and treated in three different sections: section 2 considered recent literature, section 3 discussed the results of four questionnaires conducted in top robotics forums over a period of two years and finally, section 4 summarized the security research results in robotics obtained while conducting security assessments.

Various observations were made while evaluating different sources of data which hint that the field is still immature (observations 1, 2, 3) and the practitioners are mostly yet to observe cyber-attacks (observation 4). An interesting take is that most mitigations in industry seem to be focused in the perimeter (observation 5) as this is typically the most feasible approach, however past work (Mayoral-Vilches, García-Maestro, et al., 2020) indicates that the *castle-and-moat* method won't work in robotics and instead, new approaches based on *zero-trust* should be pursued.

The lack of investment and concern that some manufacturers are showing for cybersecurity was confirmed reviewing the publicly available data on robot vulnerabilities. While some manufacturers are increasingly showing progress shortening reaction times (observation 8), others remain uncaring and persuade their users with marketing messages and strategies centered around *openness* and *community*. This appears to be a common trend (observation 7) in the collaborative robots manufactured by the danish MiR and UR, both owned by Teradyne. In light of this, additional data coming from private (non-publicly available) sources was used to evaluate the threat landscape. Based on the data, observation 9 concludes that the cooperation with security researchers and the posterior responsible disclosure leads to a reduced threat landscape (and coherently a smaller percentage of non-disclosed vulnerabilities).

The empirical results collected over a period of three years (2018-2020) performing security assessments in robots across industries indicates that most robots currently are vulnerable. The complexity of such systems is indeed one of the causes why defensive approaches are struggling to keep up with the need. Given the current maturity landscape of specialized defensive mechanisms for robots and inspired by the popular adage *"the best defense is a good offense"*, the present work recommends considering offensive practices to increase the resilience of current robotic systems. Offensive assessments can provide an external evaluation of the security readiness of a particular robot and should be integrated early within the development cycle (Mayoral-Vilches, García-Maestro, et al., 2020). Altogether, this motivates the following final observation:

Final Observation 1 *Complexity difficulties security in robotics. The inherent complexity of robotic systems leads to wide attack surfaces and a variety of potential attack vectors which manufacturers are failing to mitigate in reasonable time periods. As research advances in the field and the first commercial solutions to protect robots appear, to meet the security expectations of most immediate industries, a reverse defensive approach (an offensive one) is recommended. Periodic security assessments in collaboration with security experts will be the most effective security mechanism in the short term.*

References

Alhazmi, O., Malaiya, Y., & Ray, I. (2007). Measuring, analyzing and predicting security vulnerabilities in software systems. *Computers & Security*, *26*(3), 219 - 228. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167404806001520 doi: https://doi.org/10.1016/j.cose.2006.10.002

Alzola-Kirschgens, L., Ugarte, I. Z., Uriarte, E. G., Rosas, A. M., & Vilches, V. M. (2018). Robot hazards: from safety to security. *arXiv preprint arXiv:1806.06681*.

Bagnara, R. (2017). Misra c, for security's sake! *arXiv preprint arXiv:1705.03517*.

Balsa-Comerón, J., Guerrero-Higueras, Á. M., Rodríguez-Lera, F. J., Fernández-Llamas, C., & Matellán-Olivera, V. (2017). Cybersecurity in autonomous systems: hardening ros using encrypted communications and semantic rules. In *Iberian robotics conference* (pp. 67–78).

ISSN 2753-9997 (online). The International Journal of Cyber Forensics and Advanced Threat Investigations under a Creative Commons License, Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0)

12

Bilge, L., & Dumitraş, T. (2012). Before we knew it: An empirical study of zero-day attacks in the real world. In *Proceedings of the 2012 acm conference on computer and communications security* (pp. 833–844). New York, NY, USA: ACM. Retrieved from `http://doi.acm.org/10.1145/2382196.2382284` doi: 10.1145/2382196.2382284

Bozic, J., & Wotawa, F. (2017). Planning the attack! or how to use ai in security testing? In *Iwaise: First international workshop on artificial intelligence in security* (Vol. 50).

Breiling, B., Dieber, B., & Schartner, P. (2017, April). Secure communication for the robot operating system. In *2017 annual ieee international systems conference (syscon)* (p. 1-6). doi: 10.1109/SYSCON.2017.7934755

Caiazza, G. (2017). *Security enhancements of robot operating systems* (B.S. thesis). Università Ca'Foscari Venezia.

Caiazza, G., White, R., & Cortesi, A. (2019). Enhancing security in ros. In *Advanced computing and systems for security* (pp. 3–15). Springer.

Dieber, B., & Breiling, B. (2019). Security considerations in modular mobile manipulation. In *2019 third ieee international conference on robotic computing (irc)* (pp. 70–77).

Dieber, B., Breiling, B., Taurer, S., Kacianka, S., Rass, S., & Schartner, P. (2017, December). Security for the robot operating system. *Robot. Auton. Syst.*, *98*(C), 192–203. Retrieved from `https://doi.org/10.1016/j.robot.2017.09.017` doi: 10.1016/j.robot.2017.09.017

Dieber, B., Kacianka, S., Rass, S., & Schartner, P. (2016a). Application-level security for ros-based applications. In *Intelligent robots and systems (iros), 2016 ieee/rsj international conference on* (pp. 4477–4482).

Dieber, B., Kacianka, S., Rass, S., & Schartner, P. (2016b, Oct). Application-level security for ros-based applications. In *2016 ieee/rsj international conference on intelligent robots and systems (iros)* (p. 4477-4482). doi: 10.1109/IROS.2016.7759659

Dieber, B., Schlotzhauer, A., & Brandstötter, M. (2017). Safety & security–erfolgsfaktoren von sensitiven robotertechnologien. *e & i Elektrotechnik und Informationstechnik*, *134*(6), 299–303.

Dieber, B., White, R., Taurer, S., Breiling, B., Caiazza, G., Christensen, H., & Cortesi, A. (2020). Penetration testing ROS. In *Robot operating system (ros)* (pp. 183–225). Springer.

Finifter, M., Akhawe, D., & Wagner, D. (2013). An empirical study of vulnerability rewards programs. In *Presented as part of the 22nd USENIX security symposium (USENIX security 13)* (pp. 273–288). Washington, D.C.: USENIX. Retrieved from `https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/finifter`

Goertzel, K. M., & Feldman, L. (2009). Software survivability: where safety and security converge. In *Aiaa infotech@ aerospace conference and aiaa unmanned... unlimited conference* (p. 1922).

Guerrero-Higueras, Á. M., DeCastro-García, N., Rodríguez-Lera, F. J., & Matellán, V. (2017). Empirical analysis of cyber-attacks to an indoor real time localization system for autonomous robots. *Computers & Security*, *70*, 422–435.

Ivers, J. (2017, Mar). *Security vs. quality: What's the difference?* Retrieved from `https://www.securityweek.com/security-vs-quality-what\T1\textquoterights-difference`

Lacava, G., Marotta, A., Martinelli, F., Saracino, A., La Marra, A., Gil-Uriarte, E., & Vilches, V. M. (2020). Current research issues on cyber security in robotics.

Lera, F. J. R., Balsa, J., Casado, F., Fernández, C., Rico, F. M., & Matellán, V. (2016). Cybersecurity in autonomous systems: Evaluating the performance of hardening ros. *Málaga, Spain*, 47.

Lera, F. J. R., Llamas, C. F., Guerrero, Á. M., & Olivera, V. M. (2017). Cybersecurity of robotics and autonomous systems: Privacy and safety. In *Robotics-legal, ethical and socioeconomic impacts.* InTech.

Lera, F. J. R., Matellán, V., Balsa, J., & Casado, F. (2016). Ciberseguridad en robots autónomos: Análisis y evaluación multiplataforma del bastionado ros. *Actas Jornadas Sarteco*, 571–578.

Ma, L., Mandujano, S., Song, G., & Meunier, P. (2001). Sharing vulnerability information using a taxonomically-correct, web-based cooperative database. *Center for Education and Research in Information Assurance and Security, Purdue University*, *3*.

Mayoral-Vilches, V., Abad-Fernández, I., Pinzger, M., Rass, S., Dieber, B., Cunha, A., . . . others (2020). alurity, a toolbox for robot cybersecurity. *arXiv preprint arXiv:2010.07759*.

Mayoral-Vilches, V., Carbajo, U. A., & Gil-Uriarte, E. (2020). Industrial robot ransomware: Akerbeltz. In *2020 fourth ieee international conference on robotic computing (irc)* (pp. 432–435).

Mayoral-Vilches, V., García-Maestro, N., Towers, M., & Gil-Uriarte, E. (2020). Devsecops in robotics. *arXiv preprint arXiv:2003.10402*.

Mayoral Vilches, V., Gil-Uriarte, E., Zamalloa Ugarte, I., Olalde Mendia, G., Izquierdo Pisón, R., Alzola Kirschgens, L., . . . Cerrudo, C. (2018). Towards an open standard for assessing the severity of robot security vulnerabilities, the robot vulnerability scoring system (rvss). *arXiv preprint arXiv:1807.10357*.

Mayoral-Vilches, V., Kirschgens, L. A., Calvo, A. B., Cordero, A. H., Pisón, R. I., Vilches, D. M., . . . others (2018). Introducing the robot security framework (rsf), a standardized methodology to perform security assessments in robotics. *arXiv preprint arXiv:1806.04042*.

Mayoral-Vilches, V., Kirschgens, L. A., Gil-Uriarte, E., Hernández, A., & Dieber, B. (2018). Volatile memory forensics for the robot operating system. *arXiv preprint arXiv:1812.09492*.

Mayoral-Vilches, V., Mendia, G. O., Baskaran, X. P, Cordero, A. H., Juan, L. U. S., Gil-Uriarte, E., . . . Kirschgens, L. A. (2018). aztarna, a footprinting tool for robots. *arXiv preprint arXiv:1812.09490*.

Mayoral-Vilches, V., Pinzger, M., Rass, S., Dieber, B., & Gil-Uriarte, E. (2020). Can ros be used securely in industry? red teaming ros-industrial. *arXiv preprint arXiv:2009.08211*.

Mayoral-Vilches, V., San Juan, L. U., Carbajo, U. A., Campo, R., de Cámara, X. S., Urzelai, O., . . . Gil-Uriarte, E. (2019). Industrial robot ransomware: Akerbeltz. *arXiv preprint arXiv:1912.07714*.

Mayoral-Vilches, V., Usategui San Juan, L., Dieber, B., Ayucar Carbajo, U., & Gil-Uriarte, E. (2019). Introducing the robot vulnerability database (rvd). *arXiv e-prints*, arXiv–1912.

McClean, J., Stull, C., Farrar, C., & Mascareñas, D. (2013). A preliminary cyber-physical security assessment of the robot operating system (ros). In *Unmanned systems technology xv* (Vol. 8741, p. 874110).

McQueen, M. A., McQueen, T. A., Boyer, W. F., & Chaffin, M. R. (2009, Jan). Empirical estimates and observations of 0day vulnerabilities. In *2009 42nd hawaii international conference on system sciences* (p. 1-12). doi: 10.1109/HICSS.2009.186

Meier, L., Honegger, D., & Pollefeys, M. (2015). Px4: A node-based multithreaded open source robotics framework for deeply embedded platforms. In *2015 ieee international conference on robotics and automation (icra)* (pp. 6235–6240).

Mendia, G. O., Juan, L. U. S., Bascaran, X. P., Calvo, A. B., Cordero, A. H., Ugarte, I. Z., . . . others (2018). Robotics ctf (rctf), a playground for robot hacking. *arXiv preprint arXiv:1810.02690*.

MISRA. (2016a). *Misra c:2012 addendum 2 — coverage of misra c:2012 against iso/iec ts 17961:2013 "c secure".* (Tech. Rep.). HORIBA MIRA Limited, Nuneaton, Warwickshire, UK, April.

MISRA. (2016b). *Misra c:2012 amendment 1:"additional security guidelines for misra c: 2012,"* (Tech. Rep.). HORIBA MIRA Limited, Nuneaton, Warwickshire, UK, April.

Pfleeger, C. P., & Pfleeger, S. L. (2002). *Security in computing* (3rd ed.). Prentice Hall Professional Technical Reference.

Pichler, M., Dieber, B., & Pinzger, M. (2019). Can i depend on you? mapping the dependency and quality landscape of ros packages. In *2019 third ieee international conference on robotic computing (irc)* (pp. 78–85).

Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., . . . Ng, A. Y. (2009). Ros: an open-source robot operating system. In *Icra workshop on open source software* (Vol. 3, p. 5).

Robinson, A. (2014). *The history of robotics in manufacturing.* http://cerasis.com/2014/10/06/robotics-in-manufacturing/. (Accessed: 2018-06-05)

Rodríguez-Lera, F. J., Matellán-Olivera, V., Balsa-Comerón, J., Guerrero-Higueras, Á. M., & Fernández-Llamas, C. (2018). Message encryption in robot operating system: Collateral effects of hardening mobile robots. *Frontiers in ICT*, *5*, 2.

Shin, Y., Meneely, A., Williams, L., & Osborne, J. A. (2011, Nov). Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities. *IEEE Transactions on Software Engineering*, *37*(6), 772-787. doi: 10.1109/TSE.2010.81

Stoneburner, G. (2006, Aug). Toward a unified security-safety model. *Computer*, *39*(8), 96-97. doi: 10.1109/ MC.2006.283

Swinscow-Hall, D. (2017). *The interaction between safety and security.* `https://wwwf.imperial.ac .uk/blog/security-institute/2017/01/03/the-relationship-between-safety-and-security/`. (Accessed: 2018-05-31)

Taurer, S., Dieber, B., & Schartner, P. (2018). Secure data recording and bio-inspired functional integrity for intelligent robots. In *2018 ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 8723–8728).

Vamosi, R. (2017, Mar). *Does software quality equal software security?: Synopsys.* Retrieved from `https:// www.synopsys.com/blogs/software-security/does-software-quality-equal-software-security/`

Ward, D. D. (2006). Misra standards for automotive software.

White, R., Caiazza, G., Christensen, H., & Cortesi, A. (2019). Sros1: Using and developing secure ros1 systems. In *Robot operating system (ros)* (pp. 373–405). Springer.

White, R., Caiazza, G., Cortesi, A., Im Cho, Y., & Christensen, H. I. (2019). Black block recorder: Immutable black box logging for robots via blockchain. *IEEE Robotics and Automation Letters*, *4*(4), 3812–3819.

White, R., Caiazza, G., Jiang, C., Ou, X., Yang, Z., Cortesi, A., & Christensen, H. (2019). Network reconnaissance and vulnerability excavation of secure dds systems. In *2019 ieee european symposium on security and privacy workshops (euros&pw)* (pp. 57–66).

White, R., Christensen, D., Henrik, I., Quigley, D., et al. (2016). Sros: Securing ros over the wire, in the graph, and through the kernel. *arXiv preprint arXiv:1611.07060*.

White, R., Christensen, H. I., Caiazza, G., & Cortesi, A. (2018). Procedurally provisioned access control for robotic systems. In *2018 ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 1–9).

Young, B. (2018). *The first 'killer robot' was around back in 1979.* `https://science.howstuffworks.com/ first-killer-robot-was-around-back-in-1979.htm`.

Zheng, C., Zhang, Y., Sun, Y., & Liu, Q. (2011). Ivda: International vulnerability database alliance. In *2011 second worldwide cybersecurity summit (wcs)* (pp. 1–6).

Zhu, Q., Rass, S., Dieber, B., & Mayoral-Vilches, V. (2021). Cybersecurity in robotics: Challenges, quantitative modeling, and practice. *arXiv preprint arXiv:2103.05789*.

Appendix

A Robotics software quality, safety and security

Quality (Quality Assurance or QA for short) and **Security** are often misunderstood when it comes to software. Ivers argues (Ivers, 2017) that quality "essentially means that the software will execute according to its design and purpose" while "security means that the software will not put data or computing systems at risk of unauthorized access". Within (Ivers, 2017) one relevant question that arises is whether the quality problems are also security issues or vice versa. Ivers indicates that quality bugs can turn into security ones provided they're exploitable, and addresses the question by remarking that quality and security are critical components to a broader notion: software integrity as depicted in Figure 6.


Figure 6. According to (Ivers, 2017), Software integrity can be represented the union of both software security and quality (**Software Integrity** = Security ∪ Quality)

Coming from the same group, Vamosi (Vamosi, 2017) argues that "quality code may not always be secure, but secure code must always be quality code". This somehow conflicts with the previous view and leads one to think that secure software is a subset of quality. The author of this proposal rejects this view and argues instead that Quality and Security remain two separate properties of software that may intersect on certain aspects (e.g. testing) as depicted in Figure 7.

While the target of this research proposal is Security, Quality will also be studied given its intersection. Often, both secure and quality code share several requirements and mechanisms to assess them. This includes testing approaches such as static code testing, dynamic testing, fuzz testing or software component analysis (SCA) among others.


Figure 7. Depicts the target of this research proposal, security. Note however how security intersects with quality. This commonly refers to testing approaches such as static code testing, dynamic testing, fuzz testing or software component analysis (SCA) used both in Security and QA

In robotics there is a clear separation between Security and Quality that is best understood with scenarios involving robotic software components. For example, if one was building an industrial Autonomous Guided Vehicle (AGV) or a self-driving car, often, she/he would need to comply with coding standards (e.g. MISRA (Ward, 2006) for developing safety-critical systems). The same system's communications, however, regardless of its compliance with the coding standards, might rely on a channel that does not provide encryption or authentication and is thereby subject to eavesdropping and man-in-the-middle attacks. Security would be a
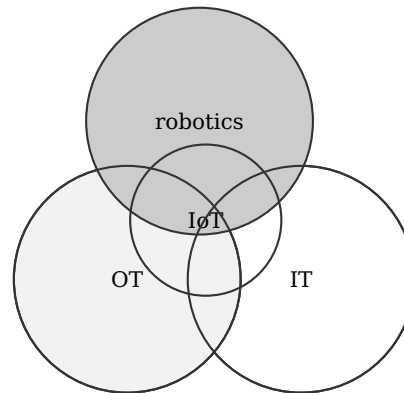
strong driver in here and as remarked by Vamosi (Ivers, 2017), "neither security nor quality would be mutually exclusive, there will be elements of both".

Quality in robotics, still on its early stages (Pichler, Dieber, & Pinzger, 2019), is often viewed as a precondition for **Safety**-critical systems. Similarly, as argued by several, safety can't be guaranteed without security (Goertzel & Feldman, 2009; Bagnara, 2017). Coding standards such as MISRA C have been extended (MISRA, 2016b, 2016a) to become the C coding standard of choice for the automotive industry and for all industries developing embedded systems that are safety-critical and/or security-critical (Bagnara, 2017). As introduced by ISO/IEC TS 17961:2013 "in practice, security-critical and safety-critical code have the same requirements". This statement is somehow supported by Goertzel (Goertzel & Feldman, 2009) but emphasized the importance of software remaining dependable under extraordinary conditions and the interconnection between safety and security in software. This same argument was later extended by Bagnara (Bagnara, 2017) who acknowledges that having embedded systems non-isolated anymore plays a key role in the relationship between safety and security. According to Bagnara, "while safety and security are distinct concepts, when it comes to connected software" (non-isolated) "not having one implies not having the other", referring to integrity.

Observation 10 *Safety and security coding (software) standards do not guarantee that the final robotic system will be secure and thereby, safe.*

In the opinion of the author of this thesis proposal, coding standards such as MISRA or ISO/IEC TS 17961:2013 for safety-critical and security-critical software components do not guarantee that the final robotic system will be secure and thereby, safe. As illustrated in the example above, robotics involves a relevant degree of system integration and inter-connectivity (non-isolated embedded systems connected together internally and potentially, externally as well). As such, both secure and ultimately safe robotics systems do not only need to ensure quality by complying against coding standards but also guarantee that they aren't exploitable by malicious attackers.

In the traditional view of *system security*, safety in often understood as "nothing bad happens naturally" while security intuitively indicates that "nothing bad happens intentionally". Acknowledging the acceptance of this view in the security community, this thesis will put special focus in the context of robotics and use a different definitions. To further understand terminology and prior art in a robotics context, Table 1 presents a summary of the concepts discussed with their interpretation applied to robotics and the corresponding sources used:

Security, as understood in Table 1 shares Integrity with Safety. As discussed in (Bagnara, 2017; Goertzel & Feldman, 2009), "the only thing that distinguishes the role of integrity in safety and security is the notion of *exceptional condition*. This reflects the fact that exception conditions are perceived as accidental (safety hazards) or intentional (security threats)". The later, security threats, are always connected to vulnerabilities. A vulnerability is a mistake in software or hardware that can be directly used by an arbitrary malicious actor or

Concept	Interpretation	Reference/s
Safety	Safety cares about the possible damage a robot may cause in its environment. Commonly used taxonomies define it as the union of integrity and the absence of hazards (Safety = Integrity + Absence of catastrophic consequences)	Alzola-Kirschgens et al. (2018); Bagnara (2017); Goertzel and Feldman (2009)
Security	Security aims at ensuring that the environment does not disturb the robot operation, also understood as that the robot will not put its data, actuators or computing systems at risk of unauthorized access. This is often summarized as Security = Confidentiality + Integrity + Availability.	Alzola-Kirschgens et al. (2018); Bagnara (2017); Goertzel and Feldman (2009); Ivers (2017)
Quality	Quality means that the robot's software will execute according to its design and purpose	Ivers (2017)
Integrity	Integrity can be described as the absence of improper (i.e., out-of-spec) system (or data) alterations under normal and exceptional conditions	Bagnara (2017)

Table 1. Summary of the concepts discussed with their interpretation applied to robotics and their references.

actress to gain access to a system or network, operating it into an undesirable manner(Pfleeger & Pfleeger, 2002). In robotics, security flaws such as vulnerabilities are of special relevance given the physical connection to the world that these systems imply. As discussed in (Alzola-Kirschgens et al., 2018), "*Safety cares about the possible damage a robot may cause in its environment, whilst security aims at ensuring that the environment does not disturb the robot operation. Safety and security are connected matters. A security-first approach is now considered as a prerequisite to ensure safe operations*". Figure 8 depicts the concepts of Safety, Quality and Security representing their relationships. The target of this research proposal remains security, however, its relationship with quality and safety must be noted.


Safety

Security | Quality

Figure 8. Pictures the relationship between safety, quality and security. In particular, safety as a super-set of security and quality. Security intersects quality in the sense that some methods are shared between both (e.g. testing). Moreover, as discussed, a safe system demands first security and quality.

Robot security vulnerabilities are potential attack points in robotic systems that can lead not only to considerable losses of data but also to safety incidents involving humans. Some claim (Zheng, Zhang, Sun, & Liu, 2011) that unresolved vulnerabilities are the main cause of loss in cyber incidents. The mitigation and patching of vulnerabilities has been an active area of research (Ma, Mandujano, Song, & Meunier, 2001; Alhazmi, Malaiya, & Ray, 2007; Shin, Meneely, Williams, & Osborne, 2011; Finifter, Akhawe, & Wagner,

2013; McQueen, McQueen, Boyer, & Chaffin, 2009; Bilge & Dumitraş, 2012) in computer science and other technological domains. Unfortunately, even with robotics being an interdisciplinary field composed from a set of heterogeneous disciplines (including computer science), to the best of the knowledge of this proposal's author and his literature review, not much vulnerability mitigation research related to robotics has been presented so far.

Observation 11 *Robot security vulnerabilities are potential attack points in robotic systems that can lead to safety incidents involving humans.*