



Faculty of Electrical Engineering

Department of Cybernetics

Programme: Electrical Engineering and Information Technology

Specialization: Artificial Intelligence and Biocybernetics

Doctoral thesis

Perirobot space representations for safe human-robot interaction

Petr Svarny

Supervisor: doc. Mgr. Matej Hoffmann, Ph.D.

Prague, 2023

Acknowledgments

First, I have to thank my supervisors Matej Hoffmann and Tomas Svoboda for the faith they had in me and that they allowed me to embark on this scientific journey, supported me during this journey, and always helped me to open new horizons. Mainly, I would like to thank Matej Hoffmann for his patience and resilience during many tight deadlines.

Additionally, my Ph.D. would not be possible without the support from our administrative staff and I would like to namely thank Petra Ivanicova, Klara Polankova, Hana Pokorna, and Kristina Lukesova.

I want to thank my colleagues for their support, inspiration, and joint work on our research. Namely these are Zdenek Straka, Jakub Rozlivek, Lukas Rustler, Jan Kristof Behrens, Michael Tesar, Hagen Lehmann, Karla Stepanova, Bedrich Himmel, Martin Sramek, and Petr Posik.

Similarly, I would like to thank my colleagues I met at the Technical University in Munich - Sami Haddadin, Saeed Abdolshah, Mazin Hamad, Erfan Shahriari, Dinmukhamed Zardykhan, Alexander Kurdas, and Robin Kirschner.

Last, but not least, I would like to thank my family – my wife, Eliska Svarna, for her support and her tolerance of my long working hours; my parents, Maria Balazsova and Pavel Svarny, for continuing to support my desire for knowledge; and my brother Michal and his family for always anchoring me with a proper life perspective.

During my PhD studies, I have been supported by the following grants and projects:

- Czech Science Foundation (GA CR) (project No. 17-15697Y),
- Czech Science Foundation (GA CR), project EXPRO (project No. 20-24186X),
- Ministry of Industry and Trade of the Czech Republic (OPPIK, project No. CZ.01.1.02/0.0/0.0/19_264/0019867),
- European Regional Development Fund, “Research Center for Informatics” (project No. CZ.02.1.01/0.0/0.0/16_019/0000765),
- Grant Agency of the Czech Technical University in Prague (grant No. SGS18/138/OHK3/2T/13),
- Grant Agency of the Czech Technical University in Prague (grant No. SGS20/128/OHK3/2T/13).

Contents

Abstract	i
Abstrakt	iii
Introduction	1
1 Structure of the thesis	1
2 Safe physical human-robot interaction	2
2.1 Human-robot collision pipeline	3
2.2 Standards for safe pHRI	3
2.3 Speed and Separation Monitoring	5
2.4 Power and Force Limiting	7
3 Efficient safe physical human-robot interaction	11
3.1 Combining SSM and PFL	11
3.2 Measuring efficiency	11
4 Perception in safe physical human-robot interaction	13
5 Contributions	15
5.1 Establish the necessity of in-situ measurements for the evaluation of ISO/TS 15066 PFL regime	15
5.2 Efficient collaboration methods for PFL and SSM regimes	15
5.3 Methodology for the deployment of collaborative robots	16
5.4 Occupancy-based representation of sensor information in the Perirobot space	16
Part I: Safety in Physical Human-Robot Interaction	19
1 3D Collision-Force-Map for Safe Human-Robot Collaboration	21
1 Collision force maps and effective mass	22
1.1 Power and Force Limiting	22
1.2 Collision-Force-Map – 2D and 3D	22
1.3 Effective mass as function of distance and height	23
1.4 Acquiring 3D Collision-Force-Map from data	23
2 Experimental setup	24
2.1 Setup and robots	24
2.2 Data collection	24
3 Experiments and results	25
3.1 3D Collision-Force-Map for UR10e	26

3.2	3D Collision-Force-Map for KUKA LBR iiwa 7R800	27
3.3	3D Collision-Force Map vs. 2D CFM vs. PFL (ISO/TS 15066)	28
3.4	Nature of dynamic impact	29
4	Discussion and conclusion	29
2	Effect of Active and Passive Soft Skins	33
1	Materials and methods	34
1.1	Robots	34
1.2	Safety stops categories	34
1.3	Collisions, their modeling, and ISO/TS 15066	35
1.4	Modeling collisions with soft protective cover	37
1.5	Transient contact simulation apparatus	38
1.6	Collision evaluation	38
1.7	AIRSKIN safety cover and collision sensor	39
1.8	Experiment setup and data collection	40
2	Results	43
2.1	Post-collision behavior	44
2.2	Impact force measurements	45
2.3	Stopping behavior effect on the impact forces	51
2.4	AIRSKIN module properties	54
3	Discussion and conclusion	55
	 Part II: Efficient Safe Collaboration	 59
3	Toward Unified PFL and SSM	61
1	Materials and methods	61
1.1	Human keypoint 3D estimation	61
1.2	Nao robot keypoints	63
1.3	Separation distance representation	63
1.4	Keypoint compensation coefficients	63
1.5	Robot control	64
1.6	HRI setup	64
2	Results	64
2.1	Basic scenario	64
2.2	Head and body discrimination	65
2.3	Dangerous tool usage	66
3	Discussion and conclusion	66
4	Combining PFL and SSM	69
1	Materials and methods	69
1.1	Robot platform	69
1.2	HRI setup	69
1.3	Software framework and robot control	70
1.4	Used human keypoints	72
1.5	Keypoint “bounding spheres”	72
1.6	Protective separation distance	74
1.7	Power and force limiting	74
1.8	Keypoint separation distance representation	75
2	Results	75

2.1	Scenario 1 and 2: Robot base vs. human keypoints	76
2.2	Scenario 3 and 4: Robot vs. human keypoints	76
2.3	Scenario 5: Addition of keypoint discrimination	76
2.4	Performance in mock task	76
3	Discussion and conclusion	77
5	Functional Mode Switching	79
1	Methodology	80
1.1	Relative velocity calculation	80
1.2	Robot functional modes	82
1.3	Smooth velocity shaping	83
2	Experiments	84
2.1	Experimental setup	85
2.2	Comparison experiment	85
2.3	Grasping experiment	85
3	Results and discussion	86
3.1	Comparison of experimental results	86
3.2	Grasping experiment results	88
3.3	Summary	88
4	Conclusion	89
	Part III: Perception and Representation	91
6	Perirobot Space	93
1	Proposed approach	93
1.1	Region of interest	94
1.2	Sensor modeling	95
1.3	Perirobot space	97
1.4	Coverage metrics	97
1.5	OctoMap modeling	97
2	Results	98
2.1	Different data interpretation experiment (Exp 1)	98
2.2	Multi-sensor coverage experiment (Exp 2)	99
2.3	Dynamic scene experiment (Exp 3)	102
3	Conclusion	105
4	Discussion	105
	Conclusions and Future Work	107
	Publications	109



Abstract

Safety is a necessity in physical human-robot collaboration, but it can come at the cost of performance. Therefore it is of interest to investigate the different ways safety can be achieved and what impact it has on performance. This dissertation investigates the space surrounding the robot—or “perirobot space”—and how the safety measures implemented in this space impact the performance. This leads to the investigation and synthesis of three aspects: the problem of safety in physical human-robot interaction, the choice of efficient methods to ascertain the safety of the interaction, and the interpretation of the perirobot space. The work shows that the currently suggested standard (ISO/TS 15066) can lead to unsafe behavior and that efficient safe human-robot collaboration asks for a complex approach. The findings are supported and demonstrated for different safety regimes (e.g., power and force limiting, speed and separation monitoring, combination) and with different robots (foremost UR10e and KUKA LBR iiwa). Finally, the thesis argues for the novel interpretation of the space surrounding robots, the so-called perirobot space.



Abstrakt

Bezpečnost je nutnou podmínkou fyzické spolupráce robota a člověka. Bohužel však v praxi dodržení bezpečnostních požadavků vede ke snížení výkonnosti celé robotické aplikace. Z toho důvodu je žádoucí zkoumat různé cesty, jakými zajistit bezpečnost a jejich vliv na výkon. Předmětem této disertace je zkoumání takových metod u kolaborativních robotů v jejich tzv. perirobotím prostoru, tj. prostoru kolem robota. Tento výzkum vede k syntéze tří oblastí: bezpečnosti ve fyzické interakci mezi člověkem a robotem, efektivity bezpečné spolupráce mezi člověkem a robotem a samotného perirobotického prostoru a jeho interpretace. Práce představuje příklady, jak je stávající návrh normy pro bezpečnou spolupráci mezi člověkem a robotem (ISO/TS 15066) nedostačující a proč je třeba přistupovat k bezpečnosti komplexněji, než jak ji prezentuje norma. Tyto závěry jsou podpořeny výsledky z již publikovaných prací autora dizertace pro různé sestavy i roboty (zejména UR10e a KUKA LBR iiwa). V závěru práce argumentuje pro nový přístup k interpretaci prostoru kolem robotů, tzv. perirobotického prostoru.

Introduction

Perception, representation, and representation of the space around the robot play a crucial role in the research fields dealing with *Physical Human-Robot Interaction* (pHRI) or *Human-Robot Collaboration* (HRC) (e.g., [1]). This space surrounding the robot—or so-called perirobot space—was never explicitly introduced although safety measures predominantly operate in this space. Classical industrial robot approaches were content with straightforward, robust safety features such as physical or virtual barriers.

Nevertheless close pHRI demands more versatile approaches. This demand leads to the introduction of collaborative robots (or so-called “cobots”), which are safe thanks to their lightweight rounded design and limited speed or payload. The expectations associated with cobots are high. They promise the empowerment of their human partners by allowing them to benefit from superior robot precision, speed, or power. Still, their uptake has been somewhat held up by their performance limitations derived from strict safety constraints and caused by the uncertainty about their proper use. This is where my dissertation thesis ties in. It studies the perirobot space and its representation for safe human-robot collaboration.

1 Structure of the thesis

Research on safe pHRI touches three main areas of robotics—safety, efficiency, and perception—that are studied often together as they strongly influence each other. In my Ph.D. thesis, I will present the state of the art for each of these areas separately, and I will also present my work from these three perspectives. First, in safety, the central focus is the definition of safe interaction and how is this safety determined. I present, especially the ISO standards-based robot safety research, in Sec. 2, and our work is discussed in detail in the Part I: Safety in Physical Human-Robot Interaction. However, safety comes often at the cost of performance. Therefore a body of research focuses also on finding solutions that are safe but also allow efficient collaboration. The summary of these approaches is in Sec. 3. My contribution to this area is then described in the Part II: Efficient Safe Collaboration. Finally, the use of various robot behaviors relies on sensors and the representation of their data. The current approaches are shown in Sec. 4 and my work is presented later in the Part III: Perception and Representation. Section 5 summarizes the contributions of this thesis.

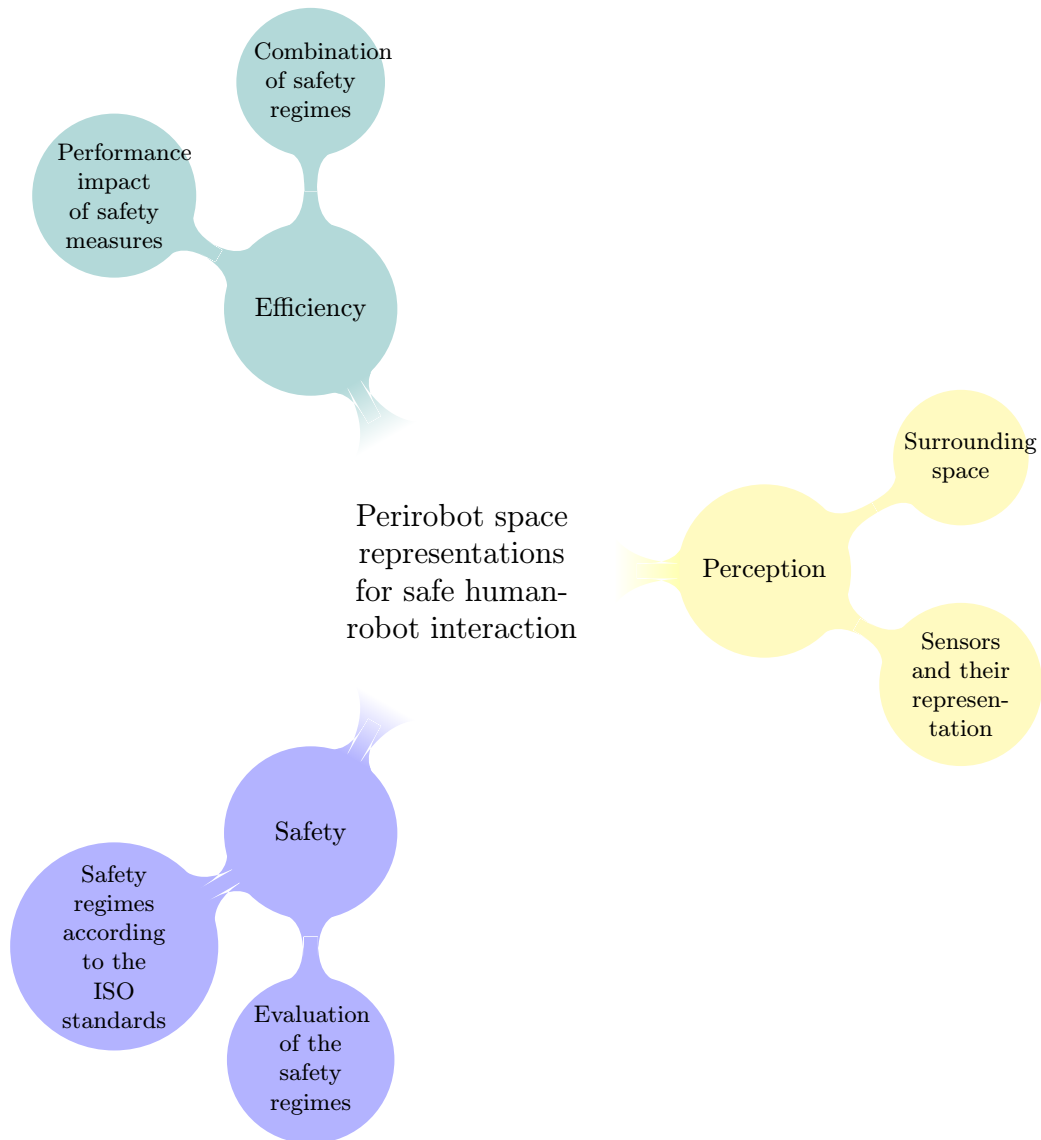


Figure 1: Topic areas and chapters of the thesis.

■ 2 Safe physical human-robot interaction

Physical Human-Robot Interaction (pHRI) and *Human-Robot Collaboration* (HRC) are dynamically growing research fields (see surveys as [1], [2], or [3] with the focus on collaboration). The so-called cobots [4] promised safe and efficient work of robots along and together with humans. They were designed with safety in mind and their control mechanisms (e.g., compliance) supported safety too.

The evaluation of whether a pHRI application is safe can be achieved in multiple ways. It can be done by using guidelines based on previous experiments or following relevant standards which should describe the best practices of the field. While safety in pHRI has been discussed and researched extensively and research thus provides a vast source of information about various aspects of safety, standards are relevant especially for industrial applications as they can have legal implications. This is true even though there exist projects such as

COVR¹ [5] that try to make collaborative robotics more attractive to the traditional industry by addressing uncertainties and suggesting safety best practices to make cobot application certification easier. Still, a lot of publications address the prescriptions suggested by these robot standards. Safety research papers even predate the inception of these standards and presents also a more general approach to understanding safety and human-robot collisions as it is not limited by the boundaries set by the standards.

Gualtieri et al. [6] present possibly the most general division of safety research into two broad areas: *prevention*, i.e., trying to avoid a collision, and *protection*, i.e., mitigating the effect of a collision [3]. These areas are not explicitly acknowledged in the standards or research, but we can notice this divide throughout the whole field. Nevertheless, a human-robot collision is not a single event. Therefore let us address the nature of human-robot collision at least briefly in order to better understand the role of protection and prevention strategies.

2.1 Human-robot collision pipeline

Haddadin [7] decomposes a human-robot collision into a series of seven phases in the so-called collision pipeline. These seven phases are:

1. Pre-collision
2. Detection
3. Isolation
4. Identification
5. Classification
6. Reaction
7. Post-collision

Therefore a different perspective can be to split safety research areas as *pre-collision*, *collision* (encompassing the five phases 2.-6. from above), and *post-collision* strategies based on the collision phases (see also [2] for a similar separation). This separation does not match exactly the previous split from Gualtieri as pre-collision strategies can also mitigate the severity of the impact but not prevent it.

2.2 Standards for safe pHRI

My thesis comments on some of the suggestions of the ISO/TS 15066 and therefore it is necessary to introduce the standards and research related to them in more detail. The appropriate standards for pHRI with collaborative robots are in decreasing order of generality ISO13849 [8] together with ISO 13850 [9], ISO 10218 [10], and ISO/TS 15066:2016 [11] (later referred to as just ISO/TS 15066). See Fig. 2 for a larger context of standards. The two most general standards, ISO 13849 and ISO 13850 cover in safety of machinery and important features that guarantee safety as the control systems or emergency stop functions, respectively. The standard ISO 10218 specifies the use of robotic devices. This standard introduces the four modes of collaboration between robots and humans. A detailed description of these modes is provided in the technical specification for collaborative operations of industrial robots ISO/TS 15066.

These standards for collaborative robotics are limited to industrial robots. Scientific investigation of safe pHRI does not need to follow these standards. It is even a common motivation of the current safety research to challenge the claims of ISO/TS 15066 (see later). Also,

¹<http://safearoundrobots.com/>

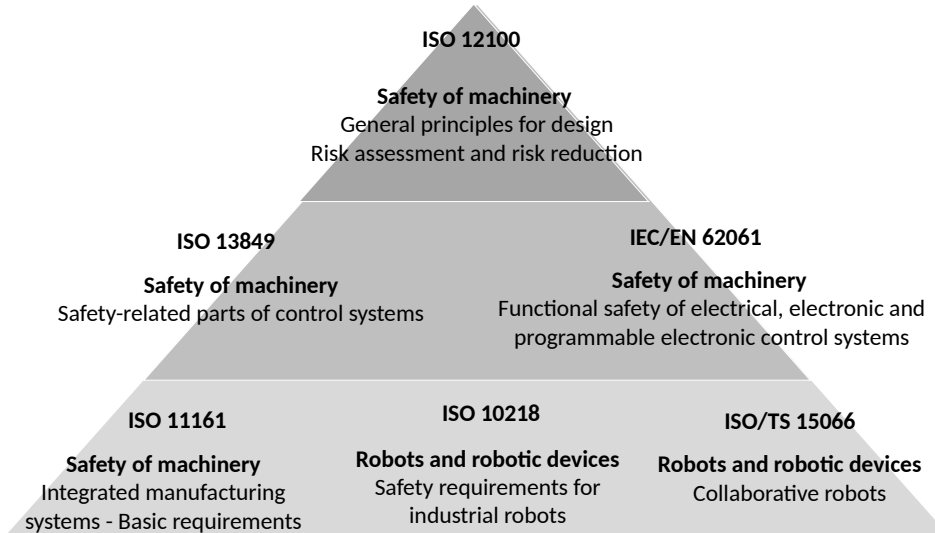


Figure 2: The pyramid representing the hierarchy of the most relevant standards for safety with industrial robots. The standard at the peak is the most general one while those at the bottom are more specific. Based on the hierarchy presented already in [12].

non-industrial robots² are not covered by the above-mentioned standards. The safety of non-industrial robots is usually satisfied by the use of lightweight design and general machinery standards [14]. For example, see the sociable humanoid robot *Pepper* with its pinch-point-less design and safe motor performance [15], which still can present risks in pHRI [16]. Some types of non-industrial robots also have their safety standards (e.g., the ISO 13482 Safety requirements for personal care robots [17], ANSI/ITSDF B56.5-2012 Safety Standard for Driverless, Automatic Guided Industrial Vehicles and Automated Functions of Manned Industrial Vehicles, and ANSI/RIA R15.06-2012 Industrial Robots and Robot Systems – Safety Requirements [18]). Interestingly though, the strict requirements of ISO/TS 15066 seem well-placed in a non-industrial context where pHRI can occur (hinted already in [14]). There are also rare cases such as the elderly-care intended GARMi robot [19] that is equipped with two Franka Panda robot arms that serve as cobots by themselves. Research done on the manipulators can be then directly applied to the GARMi robot (as shown also in the concept video [20] or our work [21]).

Interaction modes

Safety becomes a significant issue in close pHRI where dangerous collisions are possible. Humans can work together with robots under different conditions or interaction modes. We can distinguish four interaction modes. As Vicentini [22] argues, the terminology connected to interaction is used inconsistently. However, I suggest to use the following hierarchy of interaction modes based on earlier works [23, 24] and presented already in [21]. From the least interactive to the most interactive mode these would be:

Fault reaction mode. Robot is inactive, the human is has to clear the reason of the fault.

Autonomous mode. There is no shared workspace and no shared task purpose.

Coexistence mode. The human and the robot share the workspace but do not have a shared task purpose.

²I use the general term “non-industrial robots” deliberately instead of specific terms as service or care robots as the delimitation given in the standards can be problematic (see [13]).

Cooperation mode. Coexistence with a shared task purpose.

Collaboration mode. Cooperation that allows contact between the human partner and the robot.

Interaction modes present a perspective on pHRI that describes the nature of the interaction. This thesis does not focus on a particular interaction mode as some tasks ask for coexistence and others for collaboration.

Safety regimes

The ISO/TS 15066 term of collaboration does not have to necessarily mean collaboration in the above-presented sense:

3.1 collaborative operation state in which a purposely designed robot system and an operator work within a collaborative workspace
ISO/TS 15066 [11]

The “collaboration methods” used by the standards do not describe how the interaction is intended to happen but they specify the safety precautions during pHRI. As Vicentini [22] pointed out, the terminology used for task definition, i. e., interaction modes in my case, is not suitable for determining robot application safety. Due to this possible confusion between collaboration as a mode of interaction and collaboration methods as safety methods, I prefer to use the term *safety regime* or just regime for short and use collaboration to describe only the interaction mode. The ISO 10218 and ISO/TS 15066 currently list four regimes:

Safety-rated Monitored Stop (SRMS). The robot stops before the human enters the defined workspace and is kept in a monitored stop state while a person is present.

Hand-Guiding (HG). The human uses a “hand-operated” device located near the end-effector to transmit motion commands to the robot. The robot itself does not execute any independent movements and needs to be in SRMS before the operator enters the collaborative workspace.

Speed and Separation Monitoring (SSM). A collision avoidance method by monitoring the positions and velocities of both the human partner and hazardous parts of the robot. At least the minimal protective separation distance needs to be maintained at all times when the robot is moving.

Power and Force Limiting (PFL). The robot is designed or controlled so that any potential collision does not exceed the allowed collision force, pressure, and energy limits (i.e., safety thresholds) that are given for specific human body parts (e.g., hand, abdomen).

The first two regimes—SRMS and HG—are out of the scope of this thesis. Stopping the robot based on a defined workspace was used in my research. Nevertheless, it was not the target of any studies as it can be implemented by classical industrial measures and does not necessitate collaborative robots. State-of-the-art hand-guiding research mainly concerns research in control and possibly feedback from the robot if the control is achieved by moving the robot itself (see for example [25]). I present more details for the other two regimes that are used throughout the thesis extensively.

2.3 Speed and Separation Monitoring

While the Safety-Rated Monitored Stop regime clearly asks the robot to stop, the Speed and Separation Monitoring regime only asks to maintain the *minimal protective separation distance* between the human and robot while the robot is moving. This can be achieved by

stopping the robot in time, similar to SRMS. It can also lead to evasive strategies and other pre-collision phase safety measures with a finer resolution than the crude workspace approach of SRMS. Therefore we can look at the solutions for SSM as Collision Mitigation, where the robot merely stops in time, and Collision Avoidance, where the robot attempts to avoid a collision.

The requirements for a working SSM solution are :

1. sensing of the human operators' as well as robot's positions (and speeds),
2. a suitable representation of the corresponding separation distances,
3. appropriate responses of the machine (e. g., speed reduction, stop, avoidance maneuvers).

Tracking the location of the robot is possible thanks to forward kinematics and joint encoder values. The perception of the human operator's location has multiple solutions. The initial use of zone scanners presents a very crude option that permits only SRMS-like collaboration. The proliferation of two key technologies allows us to track the human operator with more detail: (i) compact and affordable RGB-D sensors (like Kinect or Realsense) and (ii) convolutional neural networks for human keypoint extraction from camera images [26,27]. With these technologies together, it is possible to detect the position of individuals and their keypoints in the collaborative workspace in real time. It is important to note that research is ahead of the industry at this point because these solutions are not yet safety-rated. The industry is moving towards 3D zone scanners (examples are Pilz SafetyEye [28] or more recently SICK safeVisionary2 [29] and the Smart Robots camera [30]) which are safety-rated but do not provide the same resolution as those used by researchers.

As mentioned, the SSM regime prohibits contact with a moving robot and this is achieved by not violating the so-called *protective separation distance*. The *protective separation distance*, S_p , is according to ISO/TS 15066 the "shortest permissible distance between any moving hazardous part of the robot system and any human in the collaborative workspace" and it is given by formula (1) in section 5.5.4.2.3 in ISO/TS 15066:

$$S_p(t_0) = S_h + S_r + S_s + C + Z_d + Z_r \quad (1)$$

with

- S_h contribution to the $S_p(t_0)$ attributable to the operator's change in location;
- S_r contribution to the $S_p(t_0)$ attributable to the robot system's reaction time;
- S_s contribution to the $S_p(t_0)$ due to the robot system's stopping distance;
- C distance that a part of the body can intrude into the sensing field before it is detected;
- Z_d position uncertainty of the operator in the collaborative workspace, as measured by the presence sensing device resulting from the sensing system measurement tolerance;
- Z_r position uncertainty of the robot system from the accuracy of the robot position measurement.

$S_p(t_0)$ can either be calculated dynamically or it can be a fixed value based on the worst-case situation. The contributions marked as S_i are determined using the robot's maximal speed v_{max} multiplied with the appropriate t_i . For example, it should be $S_r = t_r \cdot v_{max}$. Or they can be calculated, in accordance with ISO/TS 15066, as the integral: $S_i = \int_{t_0}^{T_r+T_s} v_i(t) dt$.

It is important to note that this separation distance needs to be guaranteed. Robotics researchers often create real-time obstacle avoidance solutions where the separation distance is optimized rather than guaranteed (e.g., [31,32]). Even the method of evaluation of the relative distances between the human and robot is itself a subject of discussion (see [33] or [34] for a comparison of approaches).

Research results that focused on SSM are, however, also available and flourished in the last few years. A prominent example is the work of Marvel [35,36] but also the progression of the work of Zanchettin et al. [37–40]. The latter starts with an SSM-compliant coworker automaton [41] in [37], then improve on the SSM-compliant collaboration by predicting the partner’s movements [38], by calculating the optimal avoidance path [39] or incorporating dynamic danger zones based on the robot velocity [40]. Following the above-mentioned division, we can separate the works in SSM into the two approaches of collision mitigation and collision avoidance.

Collision Mitigation

The priority of collision mitigation is to stop the robot’s movement to limit the severity of a collision and possibly even avoid it. The necessity to stop on time in dynamic environments necessitates also taking into account the velocities of the robot and the human collaborator. This leads to solutions such as “velocity obstacles” [42] or “dynamic envelope” [43]. A recent solution in this respect is the addition of the dynamic properties of the robot for accurate calculation of the necessary slowing down while avoiding collisions with an intervening human [44].

Collision Avoidance

As opposed to mere stopping, the robot can actively evade the human operator. There is a large body of work dealing with motion planning and control in dynamic environments, e.g., De Luca et al. [45] and Flacco et al. [46] deal with both pre-collision and post-collision control, Nguyen et al. [47] show a dynamic evasion approach, or also Zanchettin et al. [37] present a safe kinematic control strategy. Liu and Tomizuka [48] present a comprehensive framework that includes avoidance maneuvers of the robot and task execution while preserving safety constraints. Other examples are found in [49], where human motion tracking is integrated with optimization techniques, and in [50], where human tracking is used for fast trajectory re-planning.

2.4 Power and Force Limiting

While all HRC is “continuous, purposeful interaction associated with potential or accidental physical events” [3], only one of the above-mentioned regimes, the Power and Force Limiting regime, permits physical contact between the robot and the human while the robot is still autonomously moving, provided that the impacts stay within prescribed limits. The technical specification lists various passive or active measures (see section 5.5.5.4 in ISO/TS 15066) that limit the exerted forces (e.g., increase the contact area, use of sensing to anticipate contacts). More systematically, Haddadin and Croft [1] present two paths towards making collaboration safe: by robot design or by robot control. On the design side, the robot can have a lightweight structure, soft padding, no pinch points, and possibly elastic elements (e.g., the series elastic actuators; see [7] for a formal treatment of robots with flexible joints). Or, on the control side, collision detection and response relying on motor load measurements, force/torque, or joint torque sensing can be added or various human-aware planning and interaction planning methods can be implemented to make the application compliant with the safety requirements.

The ISO/TS 15066 prescribes limits based on the onset of pain thresholds from studies like [51]. These limits are, however, subject to a heated debate that focuses on the limits prescribed for the PFL regime and the proper evaluation of collisions, see [6, 52–57]. The force thresholds derived from the strict limits for the PFL regime enforce low operational velocities, especially if there is a risk of clamping (see for example [52] or [55]).

Also, the limits imposed by ISO/TS 15066 depend not only on the affected body part but also on the collision type. A collision itself can have two separate phases based on the force evolution during the collision [58]. An initial dynamic impact in Phase I is followed by the Phase II force profile that depends on the clamping nature of the incident (see Fig. 3). There are four possible scenarios (for details see [3] and [59]):

- unconstrained dynamic impact (no force in Phase II)
- constrained dynamic impact without clamping (diminishing force in Phase II)
- constrained dynamic impact with clamping (force is not diminishing in Phase II)
- constrained dynamic impact with oscillation (force is oscillating in Phase II)

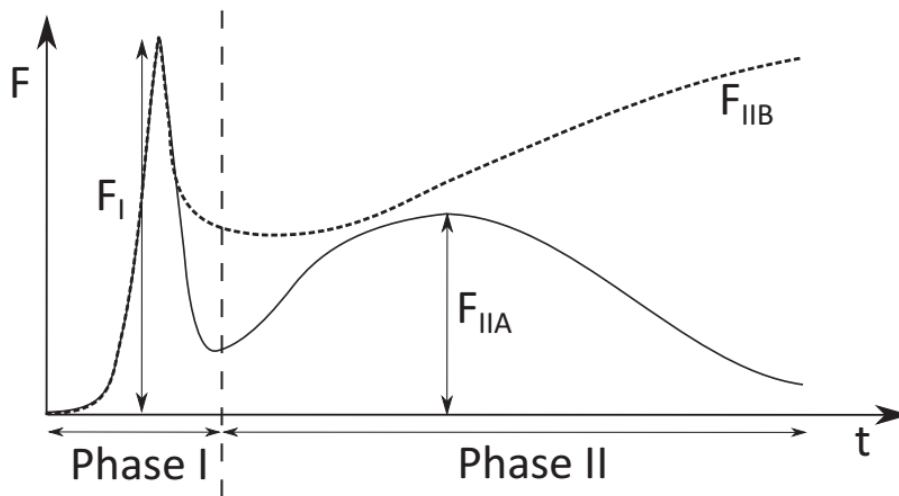


Figure 3: Collision phases from [58]. Phase I, the initial dynamic impact with the force F_I , and Phase II, either a diminishing force profile F_{IIA} in the case of no clamping or a non-diminishing force profile F_{IIB} if there is clamping.

As with SSM, different mitigation strategies need to be used for the different situations as we will see later in the presented research. In contrast to herein presented discussion about collisions, ISO/TS 15066 distinguishes only between two scenarios, a *transient* contact, i.e. dynamic impact that is unconstrained or is not followed by clamping, and *quasi-static* contact, i.e. dynamic impact followed by clamping. The limits for quasi-static contact are lower than those for transient contact.

The research approaches connected directly to PFL can be divided into three perspectives that are addressed in the subsections below:

1. What are the proper force limits for safe human-robot collisions?
2. What is a suitable impact forces model and what are its parameters?
3. What PFL deployment maximizes performance?

Establishing appropriate force limits

Especially the research of Sami Haddadin’s group tried to address force limits prescribed by the ISO/TS 15066. They studied biomechanical collision data (as in [52], [60]) to determine force limits that can lead to injury as opposed to the ISO/TS 15066 limits given by the onset of pain. Following these limits, they introduced the so-called Safe Motion Unit that would prescribe velocity limits to the robot from these biomechanical data. The original research

in [52] used injury data from mainly the automotive industry and its collision tests as in [61,62]. However, automotive industry collision data are focused on chest or whole-body collisions and not extremities. While the body region most at risk in collaborative robotics are hands. This was partially alleviated by the later focus on the injury of extremities presented in [60].

At least one other group also tried to investigate injury limits for cobots, namely, Han et al. [53] tested force limits for skin injury and came to the conclusion that only forces above 500 N lead to skin injury. Park et. al [54] conducted also a study similar to the original onset of pain studies as [51] that were the foundation of ISO/TS 15066. They explicitly mention the original motivation for the use of pressure pain thresholds instead of injury thresholds for ISO/TS 15066 as:

First, there is no accurate data on mild injury thresholds, and experiments to obtain such data have various practical limitations, including ethical issues. Second, pain thresholds are lower than mild injury thresholds.

The authors also present that age groups and body mass index can affect pain thresholds, adding another dimension to the difficulty of interpreting already subjective pain data. Additionally to the high variance between subjects, they report pain thresholds that are lower than those in ISO/TS 15066.

Therefore it is understandable that some approaches prefer to avoid the ISO/TS 15066 limits and instead try to model the collisions themselves more accurately.

Modeling impact forces

For a detailed treatment of safety aspects of human-robot collisions see [7,58]. Modeling the contacts in HRC is challenging as many parameters (mass and its distribution in colliding bodies, the behavior of robot controller upon impact, etc.) are not known. This might warrant the use of a simple approximation as the one from ISO/TS 15066. While ISO/TS 15066 presents a few formulas for the calculation of the permissible force or velocity, they use the same spring-based model. The equation A.6 from ISO/TS 15066 relates velocity (v) and maximum impact force (F_{\max}) as:

$$v \leq \frac{F_{\max}}{\sqrt{k}} \sqrt{m_R^{-1} + m_H^{-1}} = \frac{F_{\max}}{\sqrt{k \cdot \mu}}, \quad (2)$$

where m_R is the effective robot mass, m_H is the human body part mass, $\mu = (m_R^{-1} + m_H^{-1})^{-1}$ is the reduced mass of the two-body system, k is the spring constant for the human body part, and F_{\max} is the maximum impact force permitted for the given body region as established by the onset of pain studies. The mass of the robot m_R is given by the used robot, while m_H is given by the contact scenario, e. g., if collisions with the hand are considered then m_H is only the weight of the hand.

Additionally to what the standard proposes, if we investigate constrained dynamic impacts, we can approximate $m_H^{-1} \approx 0$ as in [55]. The rationale behind the approximation is the following: The impacted body part is constrained and thus immovable. Its weight in the PFL two-body spring model can therefore be considered significantly larger than the other body's, and hence approximated as infinite. This approximation allows us to simplify the situation by investigating the relative velocity as the robot velocity with the human collision region (e.g., hand) being still:

$$v \leq \frac{F_{\max}}{\sqrt{k \cdot m_R}}. \quad (3)$$

The ISO/TS 15066 has also a very simple representation of the effective robot mass m_R . It is given statically as $M/2 + m_L$, i. e., a function of the total mass of the moving parts of the robot M and the effective payload m_L of the robot.

A more sophisticated approach to calculate the effective mass was introduced already by Khatib [63]. His robot effective mass is a dynamic property given by not only the robot’s mass but also its configuration and the impact direction. It has been later adopted by many others and it is sometimes also called reflected mass (e.g., [1, 52, 55, 58, 64]). The effective mass of a manipulator in a given direction \mathbf{u} can be modeled using the formula [64]:

$$m_{\mathbf{u}}^{-1} = \mathbf{u}^T [J(\mathbf{q})M^{-1}(\mathbf{q})J^T(\mathbf{q})]\mathbf{u}, \quad (4)$$

where \mathbf{q} are the joint angles of a given position, $M(\mathbf{q})$ and $J(\mathbf{q})$ are the inertia matrix and the Jacobian matrix of the manipulator, respectively (see, e.g., [65, Ch. 3 and Ch. 7]).

Therefore it is understandable that Eq. 2 is considered by some as a too simplistic contact model (e.g., [52]), and a significant body of research tries to present alternative approaches to modeling the impact forces and thereby allow better design of the collaborative application.

Analytical approaches that try to model the impactor and determine the impact force based on them were more common in the initial years of research and even predate the introduction of ISO/TS 15066. Examples of these are [7, 66, 67] or later works as [68].

Additionally to its simplicity, the ISO/TS 15066 model takes into account only the robot weight but no other robot properties or impact factors (e. g., soft padding, impactor shape) are considered. The approaches presented in research sometimes investigate also these effects on the collision (see non-constrained blunt impacts [69] against constrained blunt impacts [70]). Shin et al. [68] also model various impactor shapes, but do not draw conclusions regarding the simple two-body model of ISO/TS 15066. A classical contact model is the Hertz contact model [71] which could account for the material properties. However, even with improved models as [72, 73], many safety-related aspects of the collision can be missed by a purely analytical model. Additionally, post-collision reactions can change the resulting exerted forces (see [1]). This favors experimental studies of collision forces that are data-driven and focus on modeling the collected collision data.

A machine learning approach is presented by Kovincic et al. [74, 75]. They suggest using collected impact data to model the impact forces using machine learning because robot reaction mechanisms play a significant role in the resulting forces and are “not known or can not be identified” [74]. A data-driven approach is also presented by Schlotzhauer et al. [76] when they introduce a 2D Collision-Force-Map. They approximate the impact forces of UR10 and UR10e robots in a pick-and-place task with a second-degree polynomial. Another example of PFL-related modeling, yet post-collision, is the use of current measuring for collision detection presented in [77].

Maximizing performance with PFL

The last avenue of research tries to stay with both the model and limits given by ISO/TS 15066 and focuses on improving the usage of PFL by itself. This often leads to the attempt to mitigate the impact effects.

One approach is to alter the robot control to limit possible impact forces. Sloth and Petersen [78] presented a method to compute safe path velocities complying with ISO/TS 15066. Other examples in [79, 80] provide a treatment of robot control taking into account the energy dissipated in possible contacts with the operator.

Another approach is to improve the robot’s impact properties with soft padding, i.e. artificial skin. Multiple variables need to be taken into account concerning skins and the

provided protection (see discussion in [22]). In addition to their passive properties, such skins can be also sensorized. These sensors can be sometimes used for control (see for example [81]), in which case they do not provide protection but merely a control option (possibly consistent with the use of the Hand Guiding regime). Properly sensorized artificial skins can also contribute to actively detecting and localize contacts and thus partake in the collision detection and isolation phases of the collision pipeline [7]. Yet, this functionality can come at a cost as Tsuji et al. [82], for example, report a trade-off between the greater thickness of the skin (i. e., improved passive properties) and proximity sensing (i. e., active properties).

■ 3 Efficient safe physical human-robot interaction

As the main drawback of collaborative robotics is the lack of performance, researchers investigated also improving performance under standard-compliant safety. While the safety regimes are presented as four separate methods in ISO/TS 15066, the specification itself also states that “collaborative operations may include one or more of the following methods”. Therefore researchers suggested various ways how to combine safe regimes to generate a more efficient collaboration. Each regime has its benefits and drawbacks and a combination of the regimes allows to alter the application’s properties to better suit the desired functionality. One example is the combination with the relatively simple safety-rated monitored stop regime, as studied in [83] and looking at various levels of interaction modes. The hand guiding regime is even closely tied to the safety-rated monitored stop regime in ISO/TS 15066 itself.

3.1 Combining SSM and PFL

The most often studied form of combination for performance increase is between SSM and PFL regimes. The SSM regime would possibly ask the robot to stop even when the human operator is in the robot’s vicinity for a longer period and thus negatively impact the overall productivity. On the other hand, the PFL regime asks the robot to move with a safe and slow velocity even when the human operator is not present. Therefore, the basic approach is to use the SSM regime’s monitoring of the human-robot distance, but instead of preventing collisions by stopping the robot, it is slowed down to the velocity considered to be safe according to PFL. The benefit of this combination is visible in Fig. 4.

The exact nature of their combination can be varied: velocity scaling based on an impact force model [68], predicting the exerted force based on motor currents [77], optimal velocity scaling [55], or the use of control barrier functions [84].

If there are multiple ways of implementing a combination of SSM and PFL, it is appropriate to ask which one of them is the most suitable and whether they are necessary. This leads to the question of measuring the efficiency or performance of HRC.

3.2 Measuring efficiency

Metrics and benchmarks are part of the safe pHRI research community. Some metrics focus on the evaluation of the quality of the provided safety or the capability to detect collisions. Examples of these are the power flux-based design metric from [73] or the reference framework for tactile robot performance from [85].

An application can as a whole be evaluated as either safe or not. However, various safety measures can impact the efficiency of the application differently. Measuring the efficiency of the application and therefore the impact of the applied safety measures is an important part of determining if those safety measures are a viable option in practice. The ISO/TS 15066 did




Figure 4: The prescribed velocity by PFL (blue) and by SSM (orange) independently for various separating distances. The dashed line shows the velocity for a combined regime of PFL and SSM. Taken from [55].

not suggest any metric for this purpose. Therefore some safety research papers introduce their metrics to establish the impact of their suggestions.

The simplest approach would be to compare the cycle time or average velocity of an application before and after the introduction of various safety measures (see for example in [86, 87]). Generally, the application with no or minimal safety measures would have the shortest cycle time and the introduction of safety measures would increase the cycle time (e.g., by stopping the robot in SSM or slowing it down with PFL). For this, either the behavior of the expected human partner should be known or various scenarios should be modeled (e.g., the effect of multiple accesses into the workspace in [87]). The cycle time approach is application-specific and therefore hard to generalize to other applications that would want to use the same safety measure. This application-specificity problem extends also to other properties of the system (like maximum velocity in [87]) or indirect metrics (like the minimum distance between the robot and human in [88]). Therefore the community attempts to introduce more general measures.

The first attempt was done by Marvel [35] when in addition to metrics evaluating the safety measures, he presented a productivity metric as:

$$P_R = \frac{\hat{\tau}}{\tau_R} \quad (5)$$

which is the ratio of the time necessary to complete the task $\hat{\tau}$ and the time necessary to complete the task with the safety measure R in effect, τ_R . The use of a ratio already permits the comparison between different tasks because it captures the change in performance relative to the task as opposed to the absolute numbers as with cycle time.

Another metric was introduced by Zanchettin et al. [37, 55]. Their robot productivity metric is based on the ratio between the end-effector velocity in the direction to the closest

point on the human body and the minimum separating distance between the human and robot, r_s :

$$r = \frac{\mathbf{n}^T \mathbf{J}_s \dot{\mathbf{q}}}{\min r_s} \quad (6)$$

Where \mathbf{n} is the unit vector for the direction from the robot point r_s to the closest point on the human body, \mathbf{J} is positional Jacobian of the point s , and $\dot{\mathbf{q}}$ is the vector of joint velocities. Higher r -metric means higher robot velocity for the same separation distance. Therefore indirectly, the r -metric is, the more efficient the collaboration is because the robot can move faster. The authors acknowledge this metric can be less intuitive to interpret, but it is not task-specific and even a single experiment provides a large number of r values and measurements for evaluation. A statistical comparison of the obtained r values is then performed to evaluate the given safety measures.

■ 4 Perception in safe physical human-robot interaction

The previously presented work in safe pHRI focused on the evaluation of safety itself or the increase of performance while maintaining safety. However, while at least two collaborative regimes presuppose some kind of workspace monitoring (namely SRMS and SSM), the above-mentioned research did not address this aspect directly. The earlier safety discussion focused merely on criteria that make the application safe but not on methods how to achieve efficient monitoring. Workspace monitoring is a crucial part of safety as space that is not monitored is unknown to the robot, and any activity in this space (or lack thereof) cannot affect the robot’s behavior. In addition to the already discussed safety, four topics are relevant to this area: sensor coverage, sensor fusion, occlusion-mitigation or gaze control, and sensor data representation.

The origin of our research is in safe human-robot interaction (HRI). While much work in safety research has been done on the side of the robot control algorithms, we focus in this article on workspace monitoring. Space that is not monitored is unknown to the robot, and any activity (or lack thereof) cannot affect the robot’s behavior. Four fields of research are relevant to this problem.

Sensor coverage. These approaches deal with covering a given area with sensor perception(see, for example, [89], [90]). Nevertheless, they often aim at coverage from a theoretical perspective (thus using only 2D coverage) and do not combine multiple types of sensors. An exception in this respect is the recently published work by Oščádal et al. [91]. They present an approach where they determine the importance of individual voxels in the shared human-robot workspace and arrange cameras to provide the highest coverage of the monitored space. However, they do not take into account occlusions caused by the human or other objects in the scene.

Sensor fusion. This area does not address proper coverage but focuses on the fusion of various sensor inputs (see [92]). Still, safety-related works such as [93, 94] leverage sensor fusion to ascertain sufficient workspace coverage. An interesting addition in this respect is [95] which evaluates two types of sensors (albeit not used together in the same setup) for 2D and 3D coverage.

Occlusion-mitigation or gaze-control. These research topics (as in [96], [97] or [98]) are also related as it shares the aim to monitor a region of space efficiently. This region, as

opposed to coverage approaches, is small (usually merely a target point). For safe interaction, all the relevant robot surrounding space needs to be considered.

Sensor data representation. The representation of the sensor readings can significantly impact the results. For example, the detected keypoints can be represented as spheres [46], capsules [48] or meshes [99] and they can be different for the robot and the human [37]. Or the relative distances between them can be also evaluated in various ways (see [33, 34] for comparisons of approaches). The approaches are often “robot-centered”. Therefore the collision primitives are centered on the robot body and possibly dynamically shaped based on the current robot velocity [31, 40, 93]. This is also true for the biologically inspired approach to “peripersonal space” representation [37, 47, 99, 100]. The safety margin is generated by a distributed array of receptive fields surrounding the electronic skin of the iCub humanoid robot. Finally, there is a large body of work dealing with motion planning and control in dynamic environments. The most recent and most related to our approach are [35, 37, 48]. While these approaches focus on robot control and not on monitoring itself, they provide a relevant basis for the representation of sensory data.

Also relevant are various studies that take into account the dynamic adjustment based on the human or robot velocity: from the kinetostatic danger field to sphere swept lines-based bounding volumes [31, 37, 40, 87, 93, 99, 101].

While the presented papers deal with camera-based monitoring (see also older survey [102]), there are other alternatives for safety-related monitoring, e.g., time-of-flight sensors [103], proximity sensing (e.g., [104], [105] or the Bosch APAS system), or distributed wireless sensor networks that track operators who do not wear any devices on themselves [106]. Especially proximity sensing presents a promising and recent avenue of research and focuses on robot-mounted sensors that can detect obstacles up to 0.5 m from the robot surface; see [107] for a thorough review.

■ 5 Contributions

The thesis presents mainly contributions in the first two areas — safety and performance. These contributions were also presented in a series of conference proceedings and a journal article summarized in Fig. 5 at the end of this section. A full list of publications is also shown in Part Publications. The specific contribution of each paper is listed in their respective chapter. The main contributions are summarized here and are the following: establish the necessity of in-situ measurements for the evaluation of the ISO/TS 15066 PFL regime, methods for efficient collaboration by combining the PFL and SSM regimes, a methodology for the deployment of collaborative robots, and the introduction of the occupancy-based unified approach to represent sensory information in the robot-surrounding space.

5.1 Establish the necessity of in-situ measurements for the evaluation of ISO/TS 15066 PFL regime

By a thorough analysis of collision data, we showed that the predicted impact forces based on the ISO/TS 15066 formula are for certain conditions too conservative and for other conditions too lenient. This contribution is presented in the Part called Safety in Physical Human-Robot Interaction.

Namely, we showed in [108] (see Chapter 1) that the formula to determine impact forces or permissible velocities does not account for the variation in impact forces in the robot workspace. We demonstrated this on two collaborative robots by conducting a series of measurements mapping the robots' impact velocities in their workspaces. We suggested the creation and use of a so-called collision force map that shows the impact forces for a specific robot in its 3D workspace. Therefore, we suggested by our comparison of the behavior of the two collaborative robots that in-situ measurements are preferable to analytical solutions to account for factors that cannot be modeled (e.g., properties of the proprietary robot controllers).

Thereafter, in [59] (see Chapter 2), we showed that soft covers can affect the impact forces in non-trivial ways and, once again, need to account for the specificity of the robot application (e.g., the robot's stopping behavior). We suggested also an addition to the ISO/TS 15066 PFL regime formula that models closer the effect of a soft cover on impact forces. Additionally, we studied the effects of various stopping behaviors on the resulting impact forces. We suggested that the collected data show the importance of the interplay between various settings and thereby support the earlier claimed necessity of in-situ measurements.

5.2 Efficient collaboration methods for PFL and SSM regimes

We presented multiple ways in which the physical human-robot collaboration can be made more efficient while still safe by a suitable combination of ISO/TS 15066 collaboration regimes, namely the Power and Force Limiting (PFL) and Speed and Separation Monitoring (SSM) regimes, see the Part Efficient Safe Collaboration.

At first, we presented the proof of concept of a simple, yet versatile, pair-wise separation monitoring approach in [109] (see Chapter 3). This approach combined the PFL regime reduced velocity with the SSM regime formulae for calculating the appropriate stopping distance and was demonstrated on a Nao robot. The pair-wise nature of the approach and its simple representation allowed for easy differentiation between contacts that necessitate a stricter approach (e.g., risk of contact with the operator's head).

This initial approach was developed and verified on a series of experiments with an industrial robot, KUKA LBR iiwa, in [110] (Chapter 4). Together with the previously mentioned

paper, these were to our knowledge the first ever published works to explicitly present the combination of PFL and SSM and its benefits.

Finally, in [21] (Chapter 5), we addressed the difference between the collaboration methods (i.e., regimes) and various interaction modes (e.g., coexistence, collaboration). We presented a novel approach where the desired interaction mode is treated as an input parameter to modulate the robot behavior. We used a smooth velocity shaping approach to ascertain that the transitions between various safety regimes impact as little as possible the robot's task performance. Additionally, we suggested that while ISO/TS 15066 proved to be overly strict in industrial settings, they can be used as a basis for non-industrial human-robot interaction.

5.3 Methodology for the deployment of collaborative robots

The conducted research was used in the creation of a methodology for the efficient and safe deployment of collaborative robots in collaboration with SEA Chomutov s.r.o. The methodology focused on the identification of risks specific to collaborative robots and their mitigation. Based on our research, the focus of the methodology was on in-situ measurements to verify the actual forces that would be exerted on the operator during a possible collision.

5.4 Occupancy-based representation of sensor information in the Perirobot space

The work with the Speed and Separation Monitoring (SSM) regime lead to the necessity of modeling the space surrounding the robot, the so-called Perirobot space. For this purpose, I presented a unified occupancy-based representation approach that allows the explicit definition of sensor data representation, the delimitation of occlusions, and clearly describes the monitored area. Additionally, this work was implemented as a publicly available project that can be used for the search for the optimal sensor setup.

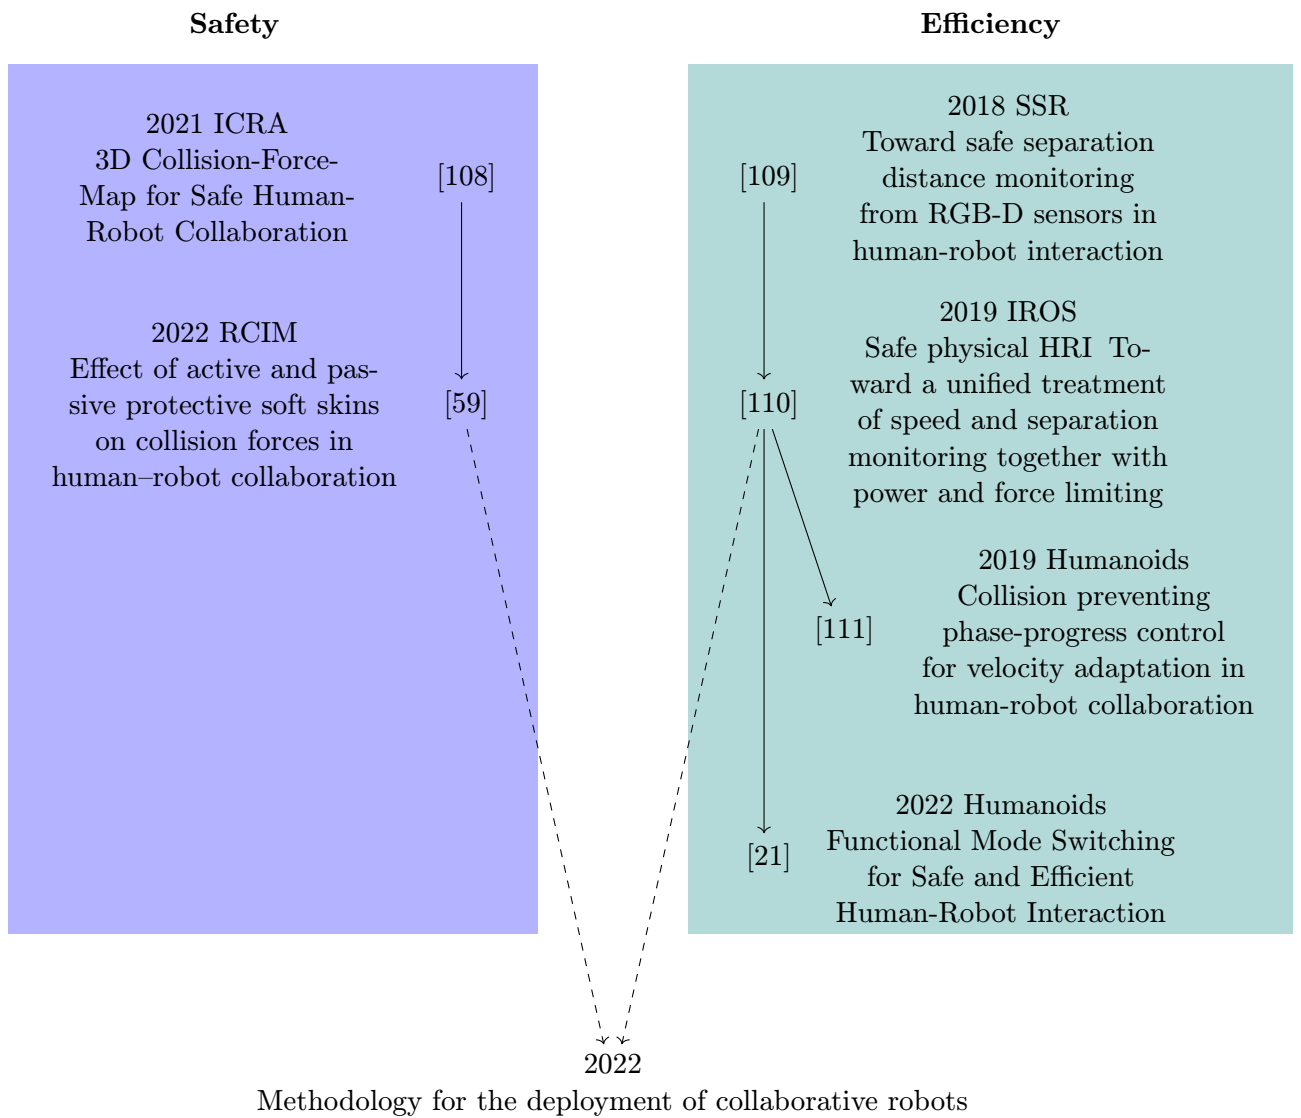


Figure 5: Thematic connections and progression of main published results with their date and publication platform. The shown conferences are International PhD Conference on Safe and Social Robotics (SSR), IEEE Conference on Robotics and Automation (ICRA), IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), and IEEE-RAS International Conference on Humanoid Robots (Humanoids). There is also one journal, Robotics and Computer-Integrated Manufacturing (RCIM).

Part I: Safety in Physical Human-Robot Interaction

This part investigates the limitations of the ISO/TS 15066 collision model and presents a large dataset modeling impact forces in the workspace of two collaborative robots in Chapter 1 and the study that investigated the effect of soft covers on the resulting impact forces in Chapter 2. I present first the common preliminaries for both of the chapters. The state of art and basic theory was presented already earlier.

Safety chapters preliminaries

Robots

We used two collaborative and one classical industrial robot for the experiments and controlled them through their standard interfaces. Each robot also has specific safety settings (see below and in Tables 2.1 and 2.2).

All robots were using the Cartesian linear movement—where the end effector follows a straight line—toward the impact. Due to various technical limitations (e.g., sensitive equipment on the KUKA iiwa’s flange), we measured collisions with the last joint’s surface.

Universal robots UR10e (UR10e) The robot has 6 degrees of freedom (DoF), weighs 33.5 kg, can carry a payload of up to 12.5 kg and has a reach of 1300 mm. The Modbus interface collected the speed and the joint states, safety regimes were collected from ROS nodes. Our UR10e is equipped with the protective skin AIRSKIN that adds extra weight (1.8 kg) to the robot, see Fig. 2.1a. If used, the skin can be connected to two different safety inputs — Emergency Stop or Safeguard Stop, see Tab. 2.1. The worst-case collection frequency for the robot speed was 800 Hz and 500 Hz for other variables.

The UR10e robots have four safety presets. We collected data with both the least restrictive (‘Pre-4’) and the second most restrictive safety (‘Pre-2’) presets.³

KUKA LBR iiwa 7 R800 (KUKA iiwa) This robot has 7 DoF, with the weight of 22.3 kg, a payload of up to 7 kg and a reach of 800 mm, see Fig. 2.1b. Two Java applications controlled the KUKA iiwa robot and collected the relevant data from the robot (1000 Hz frequency).

The KUKA iiwa assures its safety by monitoring the maximum allowed external torque with joint torque sensors in each joint. We used one external torque limit setting, namely 10 Nm or 30 Nm depending on the experiment. It could be also turned off, see the specific settings in the experiments. Three different safety stops (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’) can be triggered either by the external torque monitor or an AIRSKIN pad.

³Least restrictive preset (‘Pre-4’): Allowed power: 1000 W, Momentum: 100 kg m/s , Stopping time: 1 s, Stopping distance: 2 m, Tool speed: 5.0 m/s , Tool force: 250 N, Elbow speed: 5.0 m/s , Elbow force: 250 N. Second most restrictive (‘Pre-2’): Allowed power: 200 W, Momentum: 10 kg m/s , Stopping time: 300 ms, Stopping distance: 0.3 m, Tool speed: 0.75 m/s, Tool force: 120 N, Elbow speed: 0.75 m/s, Elbow force: 120 N.

KUKA Cybertech KR 20 R1810-2 (Cybertech) The last robot used in this chapter is a classical industrial robot with 6 DoF, weight of approximately 255 kg, a rated payload of 20 kg, and maximum reach of 1813 mm, see Fig. 2.1c. The robot was equipped with AIRSKIN module pads and controlled by a KUKA robot language program. The AIRSKIN module pad can trigger two different safety stops ('Stop 1 op', 'Stop 2'). Without sensing capabilities provided by the AIRSKIN, this robot could not be used in a collaborative operation, because the robot would not stop in case of a collision before causing harm to the human collaborator.

Measuring device

We used the *CBSF-75-Basic* impact measuring device designed for validation of collaborative applications of robots. It allows the measuring of forces in the range from 20 N up to 500 N with a certified measurement error up to 3 N. The measurement collection frequency is 1000 Hz and it starts as soon as a 20 N impact force threshold is reached and thereafter continues for 5 seconds. Therefore the pre-threshold force evolution is not collected. The peak force from the impact Phase I (see Fig. 3) was recorded and used for analysis.

In accordance with ISO/TS 15066, appropriate K1 damping materials are added to the device in order to simulate the properties of the tested human body region. Namely, we mimic impacts on the back of the non-dominant hand by using the Basic 75000 N/m device and the damping material with the hardness Sh 70 (see also [112]).

Chapter 1


3D Collision-Force-Map for Safe Human-Robot Collaboration

This chapter is based on the publication:

Svarny, P., Rozlivek, J., Rustler, L., Hoffmann, M.: 3D Collision-Force-Map for Safe Human-Robot Collaboration. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021. p. 3829-3835.

This chapter focuses on the PFL collaborative regime. We measure the forces exerted by two collaborative manipulators (UR10e and KUKA iiwa) on an impact measuring device in different positions in the robot workspace and with various velocities (Fig. 1.1). Our approach is similar to 2D CFM [76] in that we use empirical measurements and fit a function relating robot position and speed to the impact force. Newly, we establish the importance of the height, leading to a 3D Collision-Force-Map (3D CFM).¹ The dataset with the collected data is at [113].

¹An accompanying video illustrating the experiments is also available: <https://youtu.be/4eHsbe4EuHU>.



(a) UR10e.



(b) KUKA iiwa.

Figure 1.1: Setup – robots and impact measuring device. Originally published in [108].

Contributions

The contributions presented in this chapter are:

- We create a 3D collision-force-map, considering the velocity, distance from robot base, and, newly, the height in the workspace.
- We present and validate a simple data-driven model using only few samples.
- We analyze the behavior of the two manipulators upon impact, drawing important implications for their deployment in collaborative applications.

This chapter is structured into four sections: Collision force maps and effective mass (Sec. 1), Experimental setup (Sec. 2), Experiments and results (Sec. 3), and Discussion and conclusion (Sec. 4).

1 Collision force maps and effective mass

1.1 Power and Force Limiting

Some preliminaries need to be introduced additionally to the details of PFL in the Introduction, Section 2.4. We assume a mock pick and place scenario with a risk of a constrained dynamic impact on the human hand as in [76]. Contact may occur as the robot is descending towards the table, possibly clamping the hand of the operator. In practice, a risk analysis according to [114] will be required.

Because we investigate constrained dynamic impacts, we use the approximation of m_H^{-1} as discussed in the Introduction. The other variables are set based on ISO/TS 15066 as $F_{\max} = 140$ N and $k = 75000$ N/m. The moving masses of the UR and KUKA iiwa robot are approximately 30 kg and 20 kg respectively. Using the approximation from ISO/TS 15066 that the effective robot mass m_R is $M/2 + m_L$ (half of the total mass of the moving parts of the robot, plus the effective payload m_L , which is zero in our case), together with Eq. 2, would give permissible velocity up to 0.13 m/s for the UR robot and 0.16 m/s for the KUKA iiwa robot in case of clamping. If there is no clamping, the permissible force becomes 280 N and thus also the velocities are higher, namely 0.26 m/s for the UR and 0.32 m/s for the KUKA iiwa due to the weight difference between the robots.

1.2 Collision-Force-Map – 2D and 3D

The assumptions and approximations made in ISO/TS 15066 are too coarse and do not match empirical impact measurements. Schlotzhauer et al. [76] proposed a 2D Collision-Force-Map—a data-driven linear model to predict the impact force as a function of the distance from the z-axis of the robot base frame (d) and velocity (v). The model is a second degree polynomial of the form:

$$\ln(F) = \beta_0 + \beta_1 \cdot v + \beta_2 \cdot d + \beta_3 \cdot d^2 \quad (1.1)$$

The parameters are robot-, software-, and application-specific and should be found from a large number of measurements.

In this chapter, we add the height in the workspace (h) as an important additional dimension that affects the force exerted on impact. Euclidean distance in 3D between the end effector (EE) and the robot base would be a candidate representation, leading to a different 2D Collision-Force-Map. However, our empirical measurements—see Fig. 1.2—reveal a more complicated relationship between d , h , and v . For the UR10e robot (Left), the dependence of force on distance has a different profile for $h \geq 0.38$ m than for lower heights. This is true




Figure 1.2: Impact forces for different distances, heights, and velocities of EE. (Left) UR10e. (Right) KUKA iiwa 7R800 with 10 Nm external torque limit. [108]

for two different speeds. For the KUKA iiwa (Right), the contribution of height to predicting the impact forces goes down with the distance from the base.

1.3 Effective mass as function of distance and height

We sought a theoretical rationale for the observations above with effective mass (see Eq. 4 in Sec. 2.4). Although the robots have 6 (UR10e) and 7 (KUKA iiwa) degrees of freedom (DoF), the robot configurations at impact can be coarsely approximated with a 3 DoF planar manipulator. Inspired by the UR10e manipulator, we used a model with three links with masses [13, 4, 4] in kg and the length of the links [0.5, 0.45, 0.05] in m . On a grid resembling Fig. 1.4, we used the analytical solution of inverse kinematics, restricted to the “elbow up” configuration, to reach with the EE the targets on the grid—see Fig. 1.3a—and calculated the effective mass, with $u = [0, -1]$ (collision in the downward direction). We sampled the workspace more densely, giving rise to Fig. 1.3b, providing a prediction in line with Fig. 1.2. The results also suggest that the effect of d and h should be considered together and “cross-factors” are needed.

1.4 Acquiring 3D Collision-Force-Map from data

We investigated the significance of every element of the model like the one in Eq. 1.1, with additional terms in h and terms with interaction factors between d , h , and v —for 3 datasets (Tab. 1.1) simultaneously using a two-stage process. We started with the polynomial model containing all terms (variables d , v , h) up to degree three and their interaction terms up to degree three (19 terms together). We removed all terms with a p -value higher than 0.05 for all three datasets in stage one, to obtain 13 terms for stage two. In stage two, we iteratively removed terms and compared two model parameters: Root Mean Squared Error ($RMSE$) and coefficient of determination (R^2). In every iteration, the fit would typically be worse and hence $RMSE$ would increase and R^2 decrease. The term for which its removal produced the smallest change of these two parameters was removed. The change was defined as follows: $\sum_{\text{datasets}} (\Delta RMSE + 100\Delta R^2)$. The elimination procedure was stopped when this change for

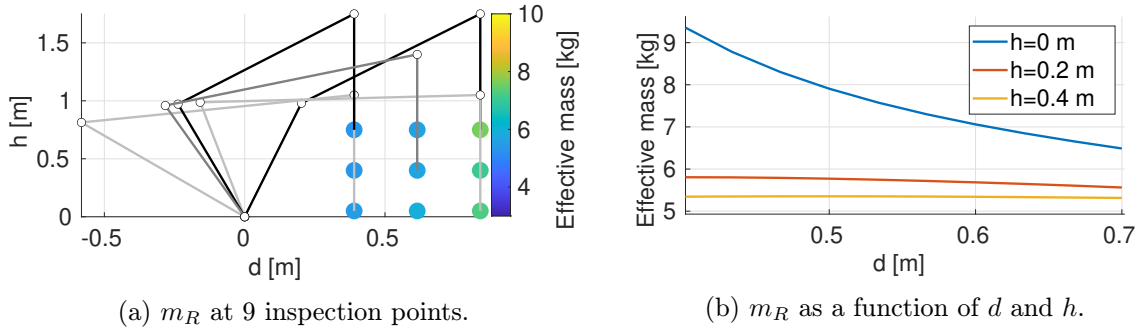


Figure 1.3: Calculating effective mass of the model 3 DoF planar manipulator with respect to distance (d) and height (h). Collision direction “down”: $u = [0, -1]$. Originally published in [108].

the term to be eliminated was bigger than 0.5 (i.e., removal of this term would make the fit significantly worse). The result of this process gave rise to Eq. 1.2:

$$\ln(F) = \beta_0 + \beta_1 \cdot v + \beta_2 \cdot d + \beta_3 \cdot d^2 + \beta_4 \cdot d \cdot h + \beta_5 \cdot h^2 + \beta_6 \cdot d^2 \cdot v + \beta_7 \cdot d \cdot v^2 + \beta_8 \cdot d \cdot h^2 \quad (1.2)$$

2 Experimental setup

2.1 Setup and robots

An overview of the setup is in Fig. 1.1². The experiments consisted of a series of impacts with the robots at different locations in the workspace and different speeds onto an impact measuring device. Both robots were commanded using the Cartesian linear movement—where the EE follows a straight line—toward the impact. As a large number of impacts were performed (more than 400 per robot in total), we preferred not to use the robot flange but the surface at the last joint instead. Robots were controlled using their standard control interfaces while experimental data were collected. We also specify the safety settings used for the experiments as they influence the robots’ overall behavior and, in particular, the response to a collision.

2.2 Data collection

Schlotzhauer et al. [76] experimentally verified the rotational symmetry assumption. Thus, a single dimension, distance from the robot base, was the only relevant parameter. In our case, it is sufficient to study a plane in the 3D workspace, varying two dimensions: d and h .

For the UR robot, the d ranged from 0.52 m to 0.89 m with increments of 0.09 m and five different heights from the level of the robot’s base starting at 0.14 m with 0.08 m increments – see Fig. 1.4 (Left). The KUKA iiwa robot has a different reach. We sampled the workspace at the following positions: d from 0.56 m to 0.86 m with an increment of 0.075 m and five heights corresponding to heights used with UR robot – see Fig. 1.4 (Right). At a given position, we performed measurements with five different velocities (0.20, 0.25, 0.30, 0.35, 0.40 m/s) in the downward direction. All measurements above the recommended limit of the impact measuring device (500 N) were discarded.

²Also visible in our video: <https://youtu.be/4eHsbe4EuHU>.

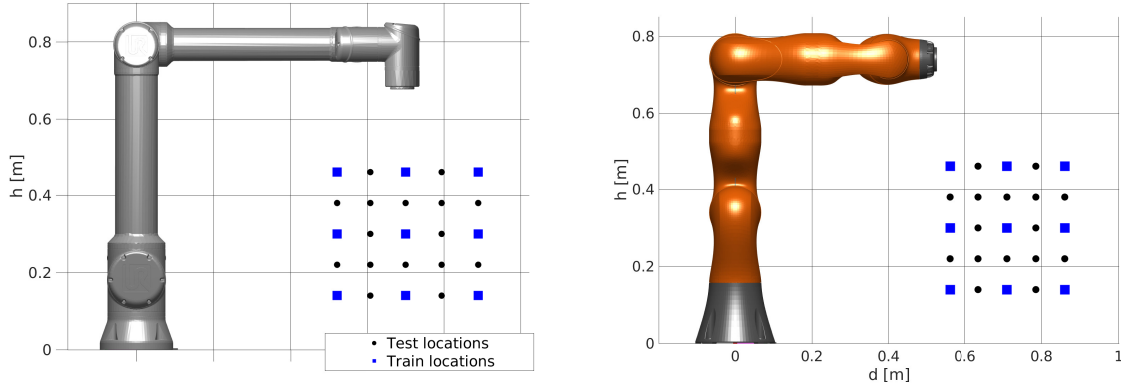


Figure 1.4: Measurements locations distribution for the two investigated robots, UR10e (left) and KUKA iiwa 7R800 (right). Originally published in [108].

dataset	samples per state	training states (used samples*)	testing states (used samples*)
UR10e	3	27 (75)	88 (249)
KUKA iiwa 30 Nm	3	27 (78)	98 (291)
KUKA iiwa 10 Nm	1	27 (26)	98 (98)

Table 1.1: Collected datasets. A “state” is a combination of distance, height, and speed. *Samples that exceeded the measuring device limit of 500 N were not used. Originally published in [108].

Training set It is our goal to develop a practical tool that can be rapidly deployed. Therefore, the number of measurements needed should be as small as possible. For training the model, we use only a subset of the grid—9 locations with blue square markers in Fig. 1.4—and 3 velocities (0.20, 0.30, 0.40 m/s). This gives rise to only 27 training measurements per robot, or 81 if every measurement is repeated 3 times. An overview is in Tab. 1.1. For the KUKA iiwa robot, the repeatability was higher. Hence, for the 10 Nm ext. torque setting, measurements were performed only once.

Testing set For every robot, 16 additional positions (black dots in Fig. 1.4) were tested with 5 velocities (0.20, 0.25, 0.30, 0.35, 0.40 m/s). For the 9 positions from the training set, only the velocities 0.25 and 0.35 m/s were added. In total, this gave 98 measurements per robot³.

Rotational symmetry verification In order to verify the assumption that rotation of the first joint does not influence the results, 117 additional measurements on the UR10e robot were performed.

■ 3 Experiments and results

Our results consist of a series of experiments in which two collaborative robots collide with an impact measuring device. First, we present and evaluate the 3D Collision-Force-Map model

³Two positions ($h = 0.3, d = 0.61$; $h = 0.3, d = 0.79$) were omitted due to the experimenter’s oversight on UR10e.




Figure 1.5: UR10e – 4D visualization of 3D CFM model predictions within the sampled robot workspace for different collision velocities. The robot base is located at $d = 0$ and $h = 0$. Black points are showing measured training data capturing the measurement grid from Fig. 1.4 (Left). Originally published in [108].

for the two robots. Second, we compare the results obtained with the 2D CFM [76] and the treatment of Power and Force Limiting in ISO/TS 15066. Finally, we present the force profiles after impact and analyze their implications.

3.1 3D Collision-Force-Map for UR10e

First, the rotational symmetry was experimentally verified using 117 measurements: 39 combinations of positions in the workspace and speed with 3 repetitions. The error, i.e. the difference in measured force on impact at different positions on the same circle (same d , h , and speed), was maximum 10 N (3.5 %), mean 1 N (0.05 %).

Second, restricting ourselves to a plane, we measured the impact forces on the grid of positions and at 5 different speeds (see Section 2.2). Every measurement was repeated 3 times, with a maximum standard deviation (SD) of these three measurements of 3.85 N and a mean of these SDs across all locations/speeds of 1.12 N. In total, 324 measurements were performed.

The training set was used to fit the model of the form in Eq. 1.2. The obtained model was:

$$\begin{aligned}
 \ln(F) = & 6.2990 + 3.3761 \cdot v - 1.1050 \cdot d - \\
 & -1.3066 \cdot d^2 - 1.5258 \cdot d \cdot h - 6.6954 \cdot h^2 + \\
 & +4.0919 \cdot d^2 \cdot v - 6.0090 \cdot d \cdot v^2 + 8.5207 \cdot d \cdot h^2
 \end{aligned} \tag{1.3}$$

Figure 1.5 shows the three variables, h , d , and v ; only the surface of this color map is visible though. As would be expected, impact forces are directly proportional to the velocity. For a fixed v and h or d , a 2D visualization is possible—green lines in Fig. 1.6.

dataset		max UE [% / N]	mean UE [% / N]	max OE [% / N]	mean OE [% / N]
UR10e	Ts	3.56 / 8.42	1.37 / 4.45	6.35 / 22.45	2.40 / 6.97
	All	3.68 / 7.07	1.50 / 4.63	6.35 / 22.45	2.16 / 6.26
KUKA iiwa 30 Nm	Ts	9.30 / 22.32	2.63 / 7.61	9.40 / 16.54	3.12 / 8.07
	All	9.30 / 22.32	2.58 / 7.37	9.40 / 16.54	3.08 / 7.93
KUKA iiwa 10 Nm	Ts	5.76 / 19.36	1.96 / 5.18	5.02 / 14.71	1.63 / 4.16
	All	5.93 / 20.20	1.94 / 5.26	5.38 / 19.48	1.59 / 4.11

Table 1.2: Accuracy of 3D CFM model with underestimation (**UE**) and overestimation (**OE**) for both the test set (**Ts**) and complete dataset (**All**). Originally published in [108].

Tab. 1.2 quantifies the accuracy of our model on the testing set (we refer to the UR10e row here). We evaluate underestimation and overestimation of the impact forces separately—the former being more critical regarding safety assessment of the application. The maximal underestimation over the testing set (see Tab. 1.1) is 3.56 % (8.42 N) and the mean underestimation is 1.37 % (4.45 N). Overestimation is higher, 6.35 % (22.45 N) at maximum and 2.40 % (6.97 N) on average. The error over the whole 3D CFM dataset is underestimating slightly more, 3.68 % (7.07 N). Overestimation is more frequent for higher force values and underestimation for lower impact forces. The higher relative overestimation is probably due to the lower density of measurements for higher forces—when impact forces surpassed 500 N.

With a bound on the underestimation, the 3D CFM can be used to determine a safe speed for a collaborative application. Adding 10% to all predicted forces—a conservative choice—and knowing d , h , and allowed impact forces, one can rearrange Eq. 1.2 to obtain the maximum safe EE velocity.

3.2 3D Collision-Force-Map for KUKA LBR iiwa 7R800

Similarly to the UR robot, restricted to a plane, we measured impact forces on the grid of positions and at 5 different speeds (Section 2.2, Fig. 1.4 (Right), Tab. 1.1 for details). We collected measurements with two different safety settings (30 and 10 Nm of max. external torque at any joint).

30 Nm external torque setting Every measurement, same location and speed, was repeated 3 times, with maximum SD of 3.09 N. The mean of these SDs across all locations/speeds was 0.58 N. The dataset is composed of 369 measurements (see Tab. 1.1). The model, 3D CFM, for this robot and settings is given by the equation:

$$\begin{aligned}
\ln(F) = & 7.0641 + 4.2943 \cdot v - 4.5286 \cdot d + \\
& + 0.9917 \cdot d^2 - 0.5795 \cdot d \cdot h - 6.0074 \cdot h^2 + \\
& + 3.9366 \cdot d^2 \cdot v - 7.2169 \cdot d \cdot v^2 + 7.0446 \cdot d \cdot h^2
\end{aligned} \tag{1.4}$$

The results for one speed (0.30 m/s) are shown in Fig. 1.6, center, with the green line—with fixed height (0.14 m, top) or distance (0.71 m, bottom).

The accuracy of the model is quantified in Tab. 1.2. The maximal underestimation error is 9.30 % (22.32 N) with a mean underestimation of 2.63 % (7.61 N) over the testing set and 2.58 % (7.37 N) over the whole 3D CFM dataset. The overestimation is comparable to the underestimation, with a mean value of 3.12 % (8.07 N) over the testing set and 3.08 % (7.93 N) over the whole 3D CFM dataset, and with the maximal error of 9.40 % (16.54 N). Both under- and overestimation are worse than in the case of the UR robot.




Figure 1.6: Impact force model comparison of 3D CFM (ours) – green, 2D CFM [76] – blue, and the value for Power and Force Limiting mode from ISO/TS 15066 – red. EE velocity = 0.30 m/s. Two approaches are presented depending on the fixed variable: (Top) Constant height of end effector in the workspace (0.14 m), (Bottom) Constant distance from base (0.70 m for UR10, 0.71 m for KUKA iiwa). Originally published in [108].

10 Nm external torque setting Due to the high repeatability of the measurements in the 30 Nm setting, deviations under the precision of the measurement device, only one measurement per position and speed was taken. The resulting dataset contains 124 measurements (see Table 1.1). The resulting model, 3D CFM, for KUKA iiwa with 10 Nm is:

$$\begin{aligned} \ln(F) = & 6.6936 + 4.9297 \cdot v - 4.4782 \cdot d + \\ & + 1.2926 \cdot d^2 - 0.3758 \cdot d \cdot h - 5.5669 \cdot h^2 + \\ & + 3.2609 \cdot d^2 \cdot v - 7.2332 \cdot d \cdot v^2 + 6.4016 \cdot d \cdot h^2. \end{aligned} \quad (1.5)$$

The results for one speed (0.30 m/s) are shown in Fig. 1.6, right, with the green line—with fixed height (0.14 m, top) or distance (0.71 m, bottom). Compared to the 30 Nm setting, the forces are on average approximately 5% lower. The accuracy of the model is quantified in Tab. 1.2. The maximal underestimation is lower than with the previous safety settings with 5.76 % (19.36 N) and an average of 1.96 % (5.18 N) over the testing set and 1.94 % (5.26 N) over the whole 3D CFM dataset. The overestimation is even lower than with the UR robot with a peak value of 5.02 % (14.71 N) and 1.63 % (4.16 N) on average over the testing set and 1.59 % (4.11 N) over the whole 3D CFM dataset.

3.3 3D Collision-Force Map vs. 2D CFM vs. PFL (ISO/TS 15066)

First, we want to compare our results with the 2D Collision-Force-Map (2D CFM) [76]. We used the least-squares method to train the 2D CFM model (Eq. 1.1) on our data, using the 0.20, 0.30, and 0.40 m/s EE velocities. A comparison for one velocity (0.30 m/s) and one height (0.14 m) is visualized in the top panels of Fig. 1.6. As the 2D CFM model does take

dataset		max UE [% / N]	mean UE [% / N]	max OE [% / N]	mean OE [% / N]
UR10e	Ts	6.51 / 14.60	2.23 / 6.33	7.56 / 11.22	1.76 / 5.17
	All	6.51 / 14.60	2.27 / 6.42	9.13 / 13.21	2.78 / 7.98
KUKA iiwa 30 Nm	Ts	12.01 / 28.83	3.94 / 10.91	7.32 / 27.67	2.73 / 7.19
	All	12.01 / 28.83	3.75 / 10.37	7.89 / 22.14	2.84 / 7.46
KUKA iiwa 10 Nm	Ts	8.70 / 29.23	3.04 / 7.59	5.24 / 21.76	2.18 / 5.82
	All	8.70 / 29.23	3.00 / 7.53	5.24 / 21.76	2.12 / 5.43

Table 1.3: Accuracy of 2D CFM models. Gray values indicate worse performance than for our 3D CFM model (Tab. 1.2). Originally published in [108].

h into account and as we have shown the forces to importantly depend on this parameter, a single 2D CFM model will fail to deliver predictions on the whole workspace. To allow for a more fair comparison, we have trained it separately for every height—blue lines in the bottom panels of Fig. 1.6. As can be seen in Tab. 1.3, the 2D CFM overestimation errors are comparable to our 3D CFM model errors (higher for UR and lower for KUKA iiwa 30 Nm dataset). On the other hand, the 2D CFM models underestimate significantly more than our 3D CFM model.

Power and Force Limiting according to ISO/TS 15066 does not take d or h into account and considers velocity only (see Section 1.1). Eq. 2 can be rearranged and F obtained. With the corresponding robot masses and $v = 0.3$, this gives rise to the red lines in Fig. 1.6. Clearly, such an approximation is insufficient. Moreover, next to overestimation, it leads also to gross underestimation of the impact forces and hence violates the safety of the human (by the very standards of ISO/TS 15066).

3.4 Nature of dynamic impact

Peak force estimation is only one component required to assess safety of a HRC application. Collision force evolution after “Phase I” (Section 1.1 and Fig. 3 in Section 2.1) is also important. Fig. 1.7 shows this for a selection of our experiments. For the UR10e robot, only Phase I is present. That is, although the scenario has a “clamping nature”, the UR10e controller makes the EE actively bounce back and thus makes the actual contact of a transient kind. On the other hand, the KUKA iiwa robot shows a prolonged damped harmonic movement upon impact.

ISO/TS 15066 prescribes maximum force thresholds for the first 0.5 s of impact (transient contact) and half this threshold afterward (quasi-static contact)—as shown in Fig. 1.7 with red dotted lines. Thus, based on our empirical findings, one could apply the higher force thresholds for the UR10e (e.g., 280 N) and only half that threshold for the KUKA iiwa (140 N), which would dramatically alter the safe speeds in the application.

4 Discussion and conclusion

Using two collaborative robots, UR10e and KUKA iiwa 7R800, we performed 934 measurements of forces exerted on the impact of the robot end-effector with an impact measuring device, with different robot velocities and at different locations in the robot workspace. The collision direction was always down, perpendicular to the table surface. We established a clear relationship between the distance from the robot base and the impact forces (in line with [76]) and, newly, also the height in the workspace—both variables being inversely proportional to




Figure 1.7: Force evolution after impact. UR10e (at $d = 0.89$ m, $h = 0.14$ m) – solid lines. KUKA iiwa (at $d = 0.86$ m, $h = 0.14$ m, 30 Nm ext. torque limit) – dashed lines. Phase I / Phase impact phase boundary (cf. Fig. 3) – cyan dotted line. Permissible force per ISO/TS 15066 – red dotted line. Originally published in [108].

the impact forces. We developed a data-driven model—3D Collision-Force-Map—that estimates the forces as a function of distance, height, and velocity, including their mutual relationships. This model is more accurate than 2D CFM [76] and PFL according to ISO/TS 15066 that does not take position in the workspace into account. Furthermore, we show that it can be trained from a limited amount of data: we sampled only 9 positions in the workspace and 3 velocities to train the model.

Thus, our main contribution is a tool that allows for rapid prototyping of a collaborative robot workspace. For quasi-static impacts on the back of the hand, a force limit of 140 N is prescribed by ISO/TS 15066, which would be based on the formula from ISO/TS 15066 limit the allowed EE speed in the whole workspace to 0.13 and 0.16 m/s for the UR10e and KUKA iiwa, respectively. Our measurements reveal that if the task is performed, for example, 0.8 m away and 0.4 m above the robot base, speeds of 0.16 m/s (UR10e) and 0.20 m/s (KUKA iiwa) stay within the prescribed force limit. Furthermore, we observe that despite the clamping nature of our scenario, the UR10e robot generates only transient contact. With the 280 N limit, 0.36 m/s will still be safe with the UR10e—an almost threefold increase. The PFL formulas from ISO/TS 15066 are insufficient—leading both to significant underestimation and overestimation at different locations in the robot workspace, and thus to unnecessarily long cycle times or even dangerous applications. The impact measuring device was firmly attached to the table. The possibility that a human operator would be moving against the robot prior to collision was thus not considered. However, we focused on quasi-static contacts where the limits are stricter. The most dangerous part of the incident is in the clamping nature.

Interestingly, the trend of the relationship between distance from / height above the robot base and the forces exerted on collision is largely consistent across two different collaborative robots and also in line with our simple 3 DoF model. However, whether the effective mass entirely determines the trends in empirically measured forces remains an open question. In addition, the impacts are not uncontrolled. The collision is detected by the robot internal controllers, generating a response, which is likely quick enough to shape the force evolution even during the first phase of the impact. Extending the effective mass models is thus impeded

by the fact that accurate inertial parameters of the manipulators are not known and the controllers are proprietary. In our case, different safety settings (external torque) resulted in different impact forces. Thus, empirical assessment of impact forces in the robot workspace seems indispensable at the moment.

It should be noted that our results are not expected to generalize to other robots or even different collision sites, directions, or kinematic configurations on the same manipulators. We concentrated on downward movement of the robot to the table, which is typical of many applications, and quasi-static contact, which represents the worst-case scenario. Impacts were made with the last robot joint, not the flange, for practical reasons. We propose an empirical method that can be deployed by robot integrators on a specific application site to quickly determine the optimal speed and position in the workspace where a task can be safely performed with maximum efficiency. The contact type and location on the robot and position in the workspace should all be set according to the application. In summary, for effective design of a collaborative application, empirical measurements are indispensable.

Chapter 2

Effect of Active and Passive Soft Skins

This chapter is based on the publication:

Svarny, P., Rozlivek, J., Rustler, L., Sramek, M., Deli, Ö., Zillich, M. and Hoffmann, M.: Effect of Active and Passive Protective Soft Skins on Collision Forces in Human–Robot Collaboration. In: *Robotics and Computer-Integrated Manufacturing*, 78. 2022.

The goal of this paper was to investigate the contribution of soft, protection-intended skin components on the overall safety of a collaborative robot system through a comprehensive study on multiple robots and with different settings. We provide a unique dataset and make it publicly available. At the same time, we propose an extension of the ISO/TS 15066 equations and demonstrate better predictions of collision forces when protective covers are used.

Contributions

The contributions presented in this chapter are:

- We perform in total 2250 collision measurements on two collaborative robots and on one traditional industrial robot, studying impact force, contact duration, clamping force, and impulse. The dataset, which can be used to develop alternative models of human-robot collisions, is publicly available¹.
- We isolate the effects of the passive padding and the active contribution of the sensor to the robot reaction.
- We systematically study the effects of additional parameters: safety stop settings, robot collision reaction, impact direction, and end effector velocity. We present insights into the interplay of these parameters and emphasize that empirical *in situ* measurements are indispensable.
- We relate the empirical measurements in different settings to the simplified prescriptions by ISO/TS 15066 and presented an extension that takes the stiffness of the protective skin and its compressible thickness into account, leading to more accurate predictions of impact forces.
- We isolate the potential of individual settings to boost productivity of a collaborative application.

The chapter is composed of three sections: Materials and methods (Sec. 1), Results (Sec. 2), and Discussion and conclusion (Sec. 3).


¹See at <https://osf.io/gwdbm>.

1 Materials and methods

In the following sections, we introduce the necessary concepts for this chapter: the safety stop categories, the characteristics of collisions, the measuring device, the setup used to simulate transient contact, the specifics of our collision evaluation approach, the AIRSKIN safety cover, and the experimental setup itself.

1.1 Robots


We used two collaborative and one classical industrial robot for the experiments and controlled them through their standard interfaces. Each robot also has specific safety settings (see below and in Tables 2.1 and 2.2). The robots themselves were discussed in more detail at the beginning of this chapter.



(a) UR10e with AIRSKIN (transient contact).



(b) KUKA iiwa with AIRSKIN module pad (quasi-static contact).



(c) KUKA Cybertech with AIRSKIN module pad (quasi-static contact).

Figure 2.1: Experiment setup — robots and the impact measuring device in both impact scenarios. Originally published in [59].

1.2 Safety stops categories

The term “Stop Category” refers to the classification of how robot motion is stopped in a safe way based on the standard IEC 60204-1 [115] and ISO 13850 [9]. There are four different types (stop category descriptions taken from [116]):

Stop Category 0. Robot motion is stopped by immediate removal of power to the robot. It is an uncontrolled stop, where the robot can deviate from the programmed path as each joint brakes as fast as possible. This protective stop is used if a safety-related limit is exceeded or in case of a fault in the safety-related parts of the control system.

Stop Category 1. Robot motion is stopped with power available to the robot to achieve the stop. Power is removed as soon as the robot stands still.

Stop Category 1 (path maintaining). Same stop category as the Stop Category 1, but the robot controller has to maintain the pre-planned task path during the controlled stop.

Stop Category 2. A controlled stop with power left available to the robot. The safety-related control system monitors that the robot stays at the stop position.

However, as illustrated in Tab. 2.1, the actual stop types of the individual robots do not exactly match those prescribed by the standard and only some of them can be triggered externally like from the protective skin. For the UR10e, Stop 0 can only be triggered through Limit violation and Fault detection and was thus not used in our experiments. The Cybertech industrial robot had only Stop 1 op and Stop 2 available. In what follows, we refer to the Emergency stop for UR10e, Stop 0 for KUKA iiwa and Stop 1 op for the Cybertech as the strictest stopping behaviors.

Stop Category	UR10e	KUKA iiwa	KUKA Cybertech
Stop 0	Limit violation Fault detection	Stop 0	Stop 0
Stop 1	—	Stop 1	Stop 1
Stop 1 (path maintaining)	Emergency stop (E-stop)	Stop 1 op	Stop 1 op
Stop 2	Safeguard stop (S-stop)	—	Stop 2

Table 2.1: Stop categories comparison between robots. Gray values indicate stops categories which can be triggered externally (i.e., AIRSKIN can trigger it). Originally published in [59].

Robot	Safety	Values		
UR10e	Preset	Pre-2	Pre-4	
	Skin	E-Stop	S-Stop	
KUKA iiwa	External torque	Stop 0	Stop 1	Stop 1 op
	Skin	Stop 0	Stop 1	Stop 1 op
KUKA Cybertech	Skin	Stop 1 op	Stop 2	

Table 2.2: Safety settings overview. The possible values for triggered stops and trigger origins or robot specific safety settings. Each robot can then have a combination of its presented settings (e.g., KUKA iiwa can have skin trigger a Stop 0 and the external torque measurement trigger a Stop 1). Originally published in [59].

1.3 Collisions, their modeling, and ISO/TS 15066

Vicentini [3] distinguishes three possible impact scenarios (see Introduction Sec. 2.4), based on our empirical study, we present three slightly different scenarios in Fig. 2.2. While Type




Figure 2.2: Collision phases and measurement times. We present three types of collisions, unconstrained dynamic impact (Type 1) and constrained dynamic impacts with clamping (Type 3), and oscillation (Type 2). The time instants on the x-axis have the following meaning: t_1 , t_2 , t_3 mark the point of peak impact force while T_1 , T_2 mark the respective dynamic collision durations. The areas J_1 , J_2 , J_3 denote the respective impulses for each type of collision. Originally published in [59].

1 represents an unconstrained dynamic impact and Type 3 matches the constrained dynamic impact with clamping, Type 2 marks a different case than was presented in the literature where oscillation occurs in Phase II. From an overall perspective, the trend of the oscillation could equate it with one of the two constrained scenarios. Nonetheless, the amplitude can make the oscillation unacceptable from the perspective of the safe force limits.

The ISO/TS 15066 presents only two scenarios, a *transient* contact, i.e., a dynamic impact that is unconstrained or it is constrained without clamping, and *quasi-static* contact, i.e., dynamic impact with clamping.

An integral part of determining the safety of the interaction is finding the appropriate safe velocities and impact forces. This is covered by the equation A.6 from ISO/TS 15066 relating velocity and impact force for transient contact:

$$v \leq \frac{F_{\max}}{\sqrt{k}} \sqrt{m_R^{-1} + m_H^{-1}} = \frac{F_{\max}}{\sqrt{k \cdot \mu}}, \quad (2.1)$$

where m_R is the effective robot mass, m_H is the human body part mass, $\mu = (m_R^{-1} + m_H^{-1})^{-1}$ is the reduced mass of the two-body system, k is the spring constant for the human body part, and F_{\max} is the maximum impact force permitted for the given body region—280 N (first 0.5 s of the impact) or 140 N (after 0.5 s of the impact) for the back of the non-dominant hand. The mass of the robot m_R is given by the used robot, while m_H is given by the scenario. For transient contact, m_H is 5.3 kg—the mass of the measuring device.

As we investigate constrained dynamic impacts, we can approximate $m_H^{-1} \approx 0$ as in [55]. This approximation allows us also to simplify the situation by investigating the relative velocity as the robot velocity with the human hand being still:

$$v \leq \frac{F_{\max}}{\sqrt{k \cdot m_R}}. \quad (2.2)$$

The ISO/TS 15066 allows considering the “effective robot mass” statically as $M/2 + m_L$, a function of the total mass of the moving parts of the robot M , and the effective payload m_L . In the case of the collaborative robots used in this work, the moving masses M of the UR10e and KUKA iiwa, are approximately 30 kg and 20 kg, respectively. These values, together with other variables set based on ISO/TS 15066 ($m_L = 0$ kg, and $k = 75000$ N/m; see Eq. 2.2) result in a permissible velocity up to 0.26 m/s for the UR10e and 0.32 m/s for KUKA iiwa if the force limit for the first 0.5 s after the collision (280 N) is considered. For the period after the first 0.5 s (clamping scenario), the stricter limit of 140 N applies, resulting in prescribing a maximum velocity of 0.13 m/s (UR10e) / 0.16 m/s (KUKA iiwa). We also use the effective mass from Khatib [63] to model the impact forces, see Introduction Sec. 2.4 for details.

1.4 Modeling collisions with soft protective cover

Eq. 2.1 is derived from the assumption that the entire kinetic energy of the robot is transferred into spring energy of the respective human body region in a fully inelastic contact (see Eq. A.2 from ISO/TS 15066):

$$E = \frac{\mu \cdot v^2}{2} = \frac{F^2}{2 \cdot k} \quad (2.3)$$

This is, however, not the case here, where the (passive or active) soft protective cover also stores spring energy. This spring energy is given by the spring constant (stiffness) of the soft protective material and the amount it can compress before being completely flattened out.

$$E = \frac{\mu \cdot v^2}{2} = \frac{F^2}{2 \cdot k} + \frac{d_s^2 \cdot k_s}{2}, \quad (2.4)$$

where d_s is the compressible thickness of the soft protective material, and k_s its spring constant. This increases the permissible speed:

$$v \leq \sqrt{\frac{F_{\max}^2}{k \cdot \mu} + \frac{d_s^2 \cdot k_s}{\mu}} \quad (2.5)$$

The stiffness of a pad can vary over its surface, e.g., an AIRSKIN module pad is harder right over the electronics housing (see also Fig. 2.15). Let us approximate the stiffness at this location with a spring constant of $k_s = 3000$ N/m and a compressible thickness of $d_s = 16$ mm. This represents the maximum compressible thickness for the AIRSKIN modules, which we are using for simplicity instead of the actual compression at the location.

For the collaborative robots equipped with AIRSKIN (see Sec. 1.7 for details) and using Eq. 2.5, this gives the following permissible velocities for quasi-static contact. During the first 0.5 s after impact (280 N force limit), the permissible velocity for the UR10e goes up from 0.26 m/s to 0.35 m/s ; for KUKA iiwa, from 0.32 to 0.43 m/s. For the period after the first 0.5 s (clamping scenario, 140 N limit), the maximum permissible velocity goes up from 0.13 to 0.26 and from 0.16 to 0.32 m/s for the UR10e and KUKA iiwa, respectively—thanks to the energy absorption by the soft protective material.

Eq. 2.5 can be rearranged to express force as follows:

$$F = \sqrt{\left(v^2 - \frac{d_s^2 \cdot k_s}{\mu}\right) \cdot (k \cdot \mu)} \quad (2.6)$$

For small velocities, it may be that $v^2 < (d_s^2 \cdot k_s)/\mu$. In these cases, no prediction is available (Fig. 2.7 and 2.10). Here, actual rather than maximum skin compression should be


considered—this is not used in this chapter. If the stiffness and compression of the protective skin is not considered, the term $(d_s^2 \cdot k_s)/\mu$ is 0 and the equation is equivalent to the corresponding ISO/TS 15066 formulas (Sec. 1.3).

This gives four variants of Eq. 2.6 which we will use later (Fig. 2.7 and 2.10): (i) ISO/TS 15066 and quasi-static contact, (ii) ISO/TS 15066 and transient contact, (iii) modified ISO/TS 15066 and quasi-static contact, and (iv) modified ISO/TS 15066 and transient contact. For quasi-static contact (i) and (iii), the simplification $m_H^{-1} = 0$ cited above is used. For (i) and (ii), the skin stiffness is not modeled; hence $(d_s^2 \cdot k_s)/\mu = 0$.


1.5 Transient contact simulation apparatus

We designed and assembled a custom mechanism to simulate *transient* contact with a frame made of aluminum profiles. The structure can be divided into two parts: a moving mass with the measuring device (see Fig. 2.3a), and a static base with ball bearing drawer slides (whole construction in Fig. 2.3b).

To secure similar conditions between individual transient experiments, we assessed the key characteristics of the apparatus. We decided to experimentally determine the smallest applied force (F_f) enough to move the moving mass horizontally and keep it the same for all experiments. Moreover, since this force affects the measured force by the measuring device, we wanted to minimize it as much as possible. Then we can assume the measured force represents a friction force and we can compute the relevant friction coefficient. For the sake of simplicity, we assume that the ball bearing drawer slides cause static friction instead of the more precise rolling resistance. The coefficient of static friction μ can be computed as $F_f = \mu_0 \cdot F_N$, where F_f is the static friction force, and F_N is the normal force. The moving mass weight (together with the measuring device) was 5.3 kg. Thus the normal force between the moving mass and the static base is approximately $F_N = 52$ N, the measured friction force is $F_f = 3.2$ N, and the coefficient of static friction is $\mu_0 = 0.062$.



(a) Moving mass



(b) Whole construction

Figure 2.3: Transient contact simulation construction. Originally published in [59].

1.6 Collision evaluation

To analyze the skin effect on the safety of the interaction, several physical quantities were investigated. Naturally, the effect on impact forces should be examined because, as mentioned before, PFL regime deals with collisions by prescribing the maximum exerted forces. Forces were compared in both phases—peak impact force in Phase I (initial dynamic impact) and

clamping force in Phase II. The skin effect on initial impact duration should also be examined, because we assume the cover prolongs the collision time. This would be the same principle as with airbags which distribute the collision impulse into a longer period of time and thus diminish the maximum force applied on the colliding human body part (head in the case of airbags) [117, Ch. 6.15].

Based on these two physical quantities, the impulse can be calculated as the integral of the measured impact force. For us, the impulse from collision start to peak impact force is practical for further use. Based on the impulse-momentum theorem (the change in momentum equals the impulse applied to it), the impulse J (in $N \cdot s$) can also be computed as:

$$J = \int_0^{t_p} F dt = \Delta p = m_r \Delta v = m_r v_0, \quad (2.7)$$

where t_p is the time of peak impact force, F is measured impact force, Δp is the change in momentum, m_r is the robot mass, Δv is difference between the initial velocity v_0 before collision and the velocity at time t_p which, we assume, is equal to zero. The impulses, peak impact forces, and relevant peak impact times of various force profiles are also shown in Fig. 2.2. As can be seen, the relation can be used to compute the expected robot mass m_r :

$$m_r = \frac{J}{v_0}, \quad (2.8)$$

where J is the computed impulse from collision start to peak impact force and v_0 is the initial velocity before collision.

For evaluation of the collision, we tracked also the reaction times, which is possible only with the UR controller. We measured the time between the onset of the collision, i.e., the measuring device starts recording as the force exceeds its 20 N threshold, and the UR controller issuing a stop status. The KUKA iiwa suspends logging when a safety response is triggered.

1.7 AIRSKIN safety cover and collision sensor

All of the robots are equipped with AIRSKIN, a soft pressure-sensitive collision sensor that covers the robot. AIRSKIN is made of individual pads, where each pad consists of an airtight hull covering a soft support structure and pressure sensors placed inside the hull [118]. Deformation of the pad as a result of contact leads to an increase in internal pressure. An increase beyond a configurable threshold (100 Pa) issues a stop signal from the sensor. Gradual pressure changes due to atmospheric changes or temperature changes of the pad are filtered out. Furthermore, each pad is slightly pressurized to 600 Pa over atmospheric pressure. Damage of the hull leads to a drop of the internal pressure, which causes the pad to go into an error state, also issuing a stop signal from the sensor.

All the electronics (2 ARM-based microcontrollers, 2 sets of internal and external MEMS pressure sensors, piezo-electric pump, and valve) are contained on the PCB mounted in each pad. All pads are connected in series, acting as opening switches for 2 safety channels, that are connected to the safety inputs of the robot controller.

These pads can either be in the form of custom shaped pads for the given robot (in our case for the UR10e robot, UR-skin in what follows; see Fig. 2.4a) or in the form of rectangular pads, AIRSKIN Module Pads (Pad in what follows; see Fig. 2.4b) used on both KUKA robots in this chapter. The UR-skin is an earlier model from 2017, while the AIRSKIN module pad is from 2019.

The recommendation from the manufacturer is to use AIRSKIN with Stop Categories 1 and 2 and not with Stop Category 0. Frequent stoppage in Stop Category 0 could damage the robot.

We wanted to verify the properties of the sensors to know if their results are comparable. We designed a measuring device that pushed the pads at various locations to study the pads' stiffness. The pushed locations were identified by a matrix — see Fig. 2.4. Due to the selected impact locations (point B2 on our matrix on both), the effect of the geometry should be low as both of the impact surfaces are flat. However, we have to take into account the surrounding area of the collision point also as the collision will happen with a flat area (the impact surface of CBSF-75-Basic) and not a single point. In addition, a large difference in the stiffness of the material could influence the measured impact force.

We used the measuring device also to determine the threshold force that activates the AIRSKIN pads. For this, the measuring device was connected with the AIRSKIN pad safety output. The device then slowly increased the exerted force upon the pad until it was stopped by the pad. The results of these measurements are presented in Sec. 2.4.




Figure 2.4: AIRSKIN pads with measurement grid. Originally published in [59].

1.8 Experiment setup and data collection

The experiments covered *quasi-static* and *transient* contacts. Each experiment consisted of a series of impacts at similar locations in the workspace of the given robot and at different velocities and different impact directions — along world frame axes.

As the robot world-frames are different, we use a right-handed coordinate frame of reference at the robot base with the y -axis along the short side of the table, the x -axis along the longer side of the table, and z -axis perpendicular to this plane — see Fig. 2.5. In this coordinate frame, our impact directions are downward ($[0, 0, -1]$), along the x -axis ($[1, 0, 0]$), and along the y -axis ($[0, 1, 0]$).

We defined 5 impact *places*—where *place* refers to a location in the workspace, impact direction, and contact type—for each robot: 3 places for quasi-static contacts (downward, along x -axis, along y -axis) and 2 places for transient contacts (along x -axis, along y -axis). Their world coordinates can be found in Tab. 2.4. The locations of every place for a given direction are similar for both types of contacts, making a comparison possible.

Robot & setup	Contact type	Skin type	Skin settings	Safety settings	Velocities [m/s]	Samples	
UR10e	quasi-static	UR-skin	S-stop E-stop Passive	Pre-4	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5	252	
		-	No skin				
	quasi-static	UR-skin	S-stop Passive	Pre-2	0.2, 0.25, 0.3 0.35, 0.4	135	
		-	No skin				
	transient	UR-skin	S-stop Passive	Pre-4	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5, 0.6, 0.7	162	
		-	No skin				
KUKA iiwa	quasi-static	Pad	Stop 0 Stop 1 Stop 1 op Passive	Stop 0	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5	315	
		-	No skin				
	quasi-static	Pad	Stop 0 Stop 1 Stop 1 op Passive	Stop 1	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5	268	
		-	No skin				
	quasi-static	Pad	Stop 0 Stop 1 Stop 1 op Passive	Stop 1 op	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5	293	
		-	No				
	quasi-static	Pad	Stop 0 Stop 1 Stop 1 op	Off	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5	173	
		-	No skin				
	transient	Pad	Stop 0 Passive	Stop 0	0.2, 0.25, 0.3 0.35, 0.4, 0.45 0.5, 0.6, 0.7	162	
		-	No skin				
	UR10e Pads comparison	quasi-static	UR-skin	S-stop Passive	Pre-4	0.2, 0.3, 0.4	135
			Pad	S-stop Passive			
-			No skin				
transient		UR-skin	S-stop Passive	Pre-4	0.2, 0.4, 0.6	90	
		Pad	S-stop Passive				
		-	No skin				
KUKA Cybertech	quasi-static	Pad	Stop 1 op Stop 2	—	0.2, 0.25, 0.3 0.35, 0.4	20	
	transient	Pad	Stop 1 op	—	0.2, 0.3, 0.4, 0.5, 0.6	20	
-	-	No skin					


Table 2.3: Datasets overview. For details see Sec. 1.2, 1.7, and 1.8. Originally published in [59].

We distinguish three principal cases of experiments:


1. **no skin** – no protective cover attached to the collision site
2. **passive skin** – protective cover attached to collision site, pressurized but without any collision detection and reaction (on the part of the skin cover)
3. **active skin** – protective cover at collision site with collision detection and connection to robot controller

The last case can be further divided based on the different safety settings for the skin (e.g., Emergency stop and Safeguard stop for UR10e robot). Data were separated into individual datasets. Their overview can be found in Tab. 2.3 and they are publicly available². Every measurement on the UR10e and KUKA iiwa robots was repeated 3 times. On the KUKA Cybertech, in order to limit mechanical stress to the heavy robot, single measurements were taken.


²The data location is <https://osf.io/gwdbm>.



(a) Impact direction downward $([0, 0, -1])$, quasi-static case



(b) Impact direction along the y -axis $([0, 1, 0])$, transient case



(c) Impact direction along the x -axis $([1, 0, 0])$, transient case

Figure 2.5: Impact directions and represented cases. The origin of the world reference frame is always at the base frame of the robot. The coordinates in the image are moved from the base frame for visibility. Originally published in [59].

UR10e For the UR10e robot, we divided the collected data into three different datasets. The first dataset contains *transient* contact impacts for active (Safeguard stop), passive, and no skin at 9 different velocities (0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.6, and 0.7 m/s) with the least restrictive safety preset (‘Pre-4’). *Quasi-static* contact impacts were collected for active (Safeguard stop), active (Emergency stop), passive, and no skin at 7 different velocities (from 0.2 to 0.5 m/s with 0.05 m/s increment) with the least restrictive safety preset (‘Pre-4’). And the third dataset consists of *quasi-static* contact impacts for active (Safeguard stop), passive, and no skin in 5 different velocities (from 0.2 to 0.4 m/s with increment 0.05 m/s) with the second most restrictive safety preset (‘Pre-2’).

KUKA iiwa Five datasets were collected with the KUKA iiwa robot. The first dataset contains *transient* contact impacts for active (‘Stop 0’), passive, and no skin in 9 different velocities (0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.6, and 0.7 m/s), with Stop 0 for external torque limit 10 N.

The remaining datasets consist of *quasi-static* contact impacts. Three of them contain impacts with all five skin settings combinations — active skin (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’), passive skin and no skin — at 7 different velocities (from 0.2 to 0.5 m/s with 0.05 m/s increment). Each of these datasets also has a different safety stop (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’) setting for external torque limit 10 N. The last dataset, with no safety stop for external torque in 7 different velocities (from 0.2 to 0.5 m/s with increment 0.05 m/s), has only active skin safety settings combinations (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’). Impacts for the combination with no safety stop for neither external torque, nor skin exceeded 500 N even for low velocities.

Place			Coordinates [m]		
Type	#	direction vector	UR10e	KUKA iiwa	KUKA Cybertech
quasi-static	0	$\begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix}$	0.85	0.66	
			0.27	0.00	–
			0.14	0.14	
	1	$\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$	0.79	0.35	0.00
			0.14	0.14	0.90
			0.16	0.16	0.18
2	$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$	0.80	0.37	0.25	
		–0.22	–0.31	0.75	
		0.16	0.16	0.18	
transient	3	$\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$	0.79	0.35	0.00
			0.18	0.10	0.90
			0.16	0.16	0.18
	4	$\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$	0.82	0.33	0.25
			–0.22	–0.31	0.75
			0.16	0.16	0.18

Table 2.4: World frame coordinates for the impact locations. The number ‘#’ identifies the place, the direction vector is given in the world frame and the coordinates are given in the world frame. The origin of the world frame is located at the base frame of the robot. Originally published in [59].

UR10e pads comparison Since a different AIRSKIN pad for each robot was used, we decided to set up another experiment with the UR10e robot having the impact point covered either by the Pad or UR-skin (experiment “UR10e pads comparison” in Tab. 2.3). In this case, the data are divided into two datasets. The first dataset contains *transient* contact impacts for active Pad or UR-skin (Safeguard stop), passive Pad or UR-skin, and no skin at 3 different velocities (0.2, 0.4, and 0.6 m/s) with the least restrictive safety preset (‘Pre-4’). The second one consists of *quasi-static* contact impacts for active Pad or UR-skin (Safeguard stop), passive Pad or UR-skin, and no skin in 3 different velocities (0.2, 0.3, and 0.4 m/s) with the least restrictive safety preset (‘Pre-4’).

KUKA KR Cybertech Since KUKA Cybertech is a non-collaborative robot with a mass of around 250 kg, quasi-static collisions without skin or with passive skin would be dangerous (for the robot, AIRSKIN, and the measuring device). For that reason, we collected a *quasi-static* dataset only for the active skin in both externally triggered stop categories (‘Stop 1 op’, ‘Stop 2’) at 5 different velocities (0.2, 0.25, 0.3, 0.35, 0.4 m/s) in directions along x -axis and y -axis (data along z -axis were not collected). The *transient* contact impacts were collected for active skin (‘Stop 1 op’) and no skin at 5 different velocities (0.2, 0.3, 0.4, 0.5, 0.6 m/s), because there was no difference in resulting forces between stop categories and active/passive skin for the transient experiments.

2 Results

The presented results consist of four separate parts. First, we present the post-collision behaviors. This is followed by the impact force measurements of the studied robots (UR10e, KUKA iiwa, and Cybertech) in various combinations of AIRSKIN and robot stop settings (see Tab. 2.3). Then, the effects of the stopping behavior on the impact force are presented. Last, we present the measurements of the AIRSKIN’s properties, stiffness and activation threshold

force, to demonstrate the possible effect of the skin’s properties.³

2.1 Post-collision behavior

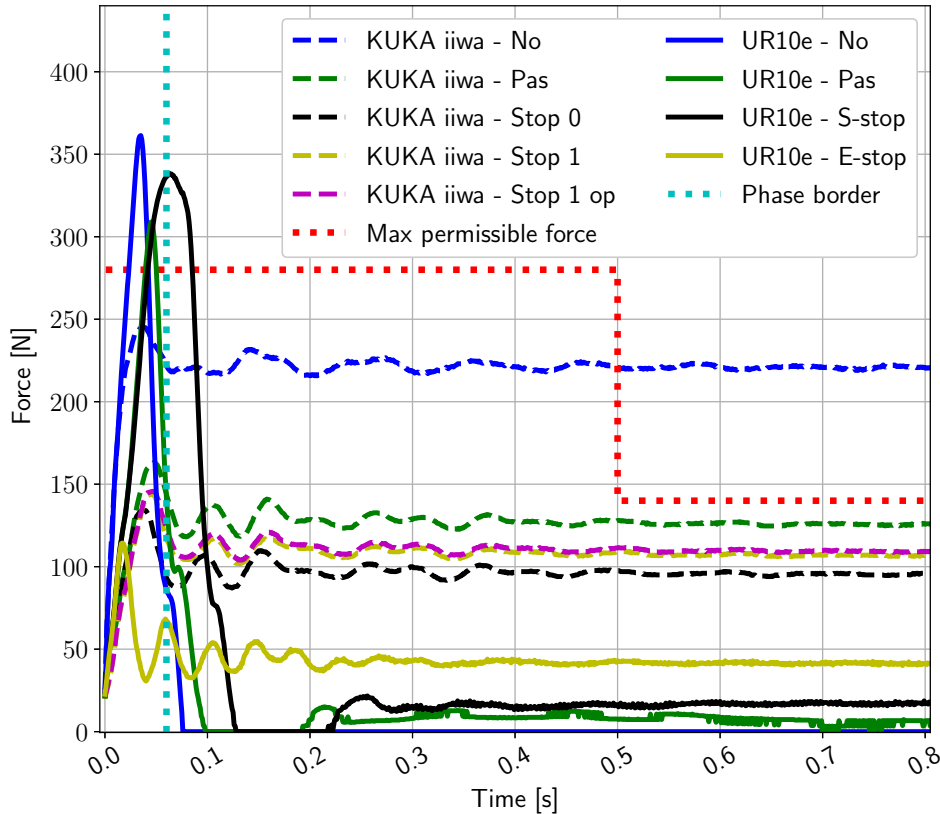


Figure 2.6: Force evolution after impact for the velocity 0.3 m/s at Place 2. UR10e – solid lines. KUKA iiwa – dashed lines. The AIRSKIN module is either absent (‘No’), active (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’, ‘S-stop’, ‘E-stop’) or merely pressurized but not active (‘Pas’) (see Tab. 2.1). An example of an actual phase I / impact phase boundary (cf. Tab. 2.2), namely for KUKA - skin No — cyan dotted vertical line. Permissible force per ISO/TS 15066— red dotted horizontal line. Originally published in [59].

The robots, due to their proprietary controllers, present different reaction behavior (see also [108]). Upon quasi-static impact, the UR10e generally bounced back, while the KUKA robots stayed at the impact position. This means the impact duration needs to be interpreted differently between the robots and even between collisions. A Safeguard stop based UR10e impact is clearly delimited in the measurements as the robot stops exerting force upon the measuring device. The KUKA iiwa, however, shows a prolonged damped harmonic movement upon impact. See Fig. 2.6 for comparison between the collaborative robots. In addition, we observed that the Emergency stop of the UR10e also leads to this damped harmonic movement. The probable cause of this movement is the oscillation of the robot and its joint motors after the abrupt stoppage. In these cases, we delimited the impact duration by the measurement onset and the first minimum after the impact peak. The Cybertech collision force profile did not allow the distinction of this first minimum, as the robot continued to

³A video illustrating a selection of the experiments is available at our Youtube channel: https://youtu.be/yqEjnK9_hCg.

push against the measuring device (see Fig. 2.11). Therefore some measurements (e.g., phase I collision duration) cannot be presented. However, let us add that the use of AIRSKIN changes the magnitude of the forces and not the shape of the force profile. The force profile depends on the robot and its respective stopping behaviors (see the differences in Fig. 2.6).

2.2 Impact force measurements

We present the collected impact force measurements hereafter. The three robots employed differ in their mechanical properties (mass, degrees of freedom etc.), controllers, and safety settings. Hence, the results across the robots should not be directly compared. Therefore we focused on the comparison between different settings for a given robot (active vs. passive vs. no skin, safety settings, robot presets). This allowed us to look at common trends between the robots.

UR10e

We present the results in Fig. 2.7. The figure can be investigated row by row as each row is a different perspective on the collision. A distinction is visible between the behavior of the first three columns representing quasi-static collisions and the last two columns, the transient cases.

The impact forces grow with the velocity in a similar manner independent of the applied protective measures. However, it is visible that the use of the strict Emergency stop with Pre-4 ('skin E-stop'), leads to lower impact forces than any other method. It also resulted in impact forces in general below the allowed force threshold. Other settings are not so clearly delimited. For example, the active use of the skin (markers filled only on the left side) generally yielded lower impact forces than passive skin (markers filled only on the right). However, this is not true for Place 2 and Pre-4, where the active skin ('S-stop') has higher impact force values than passive skin ('Pas'). Nevertheless, the first row shows that active skin leads consistently to lower impact forces than when there was no skin ('No Pre-2' and 'No Pre-4'). In transient collisions, passive skin leads to the same impact force as active skin; see the matching semicircles in the figure.

The dashed lines illustrate force predictions with Eq. 2.6 using the standard version from ISO/TS 15066 (Sec. 1.3) (blue) where no compliant cover is considered and using the modification taking the stiffness and compressible thickness of the cover into account when AIRSKIN is used (green). The modified model prediction (available only for bigger velocities; see Sec. 1.4 for details) indeed better estimates the impact forces when protective skin is employed.

Contact duration in the quasi-static cases shows predominantly three trends. First, for the no skin or E-stop case, the duration is independent of the velocity and almost constant. Second, for S-stop, the contact duration gets shorter with higher velocity (see S-stop). The last trend is transient contacts. These show a very short contact duration for all the settings with a slightly longer duration when AIRSKIN was used. This is probably due to its softness and deformation upon collision. Especially in the case of the passive skin, the absorption of the collision by the skin delays the collision detection by the robot by delaying the moment when the critical threshold is exceeded.

The clamping force played a significant role only when E-stop was used. Thus while this strict stopping behavior diminished the impact force, it also lead to clamping because the robot joints were stopped by a path maintaining Stop 1 and could not bounce of. Nevertheless, the clamping forces were safely lower than the maximum permissible clamping force (140 N).

The impulse was overall higher or comparable between the use of the skin and no skin. Lowest impulses were measured again with the use of the E-stop. Depending on the place,

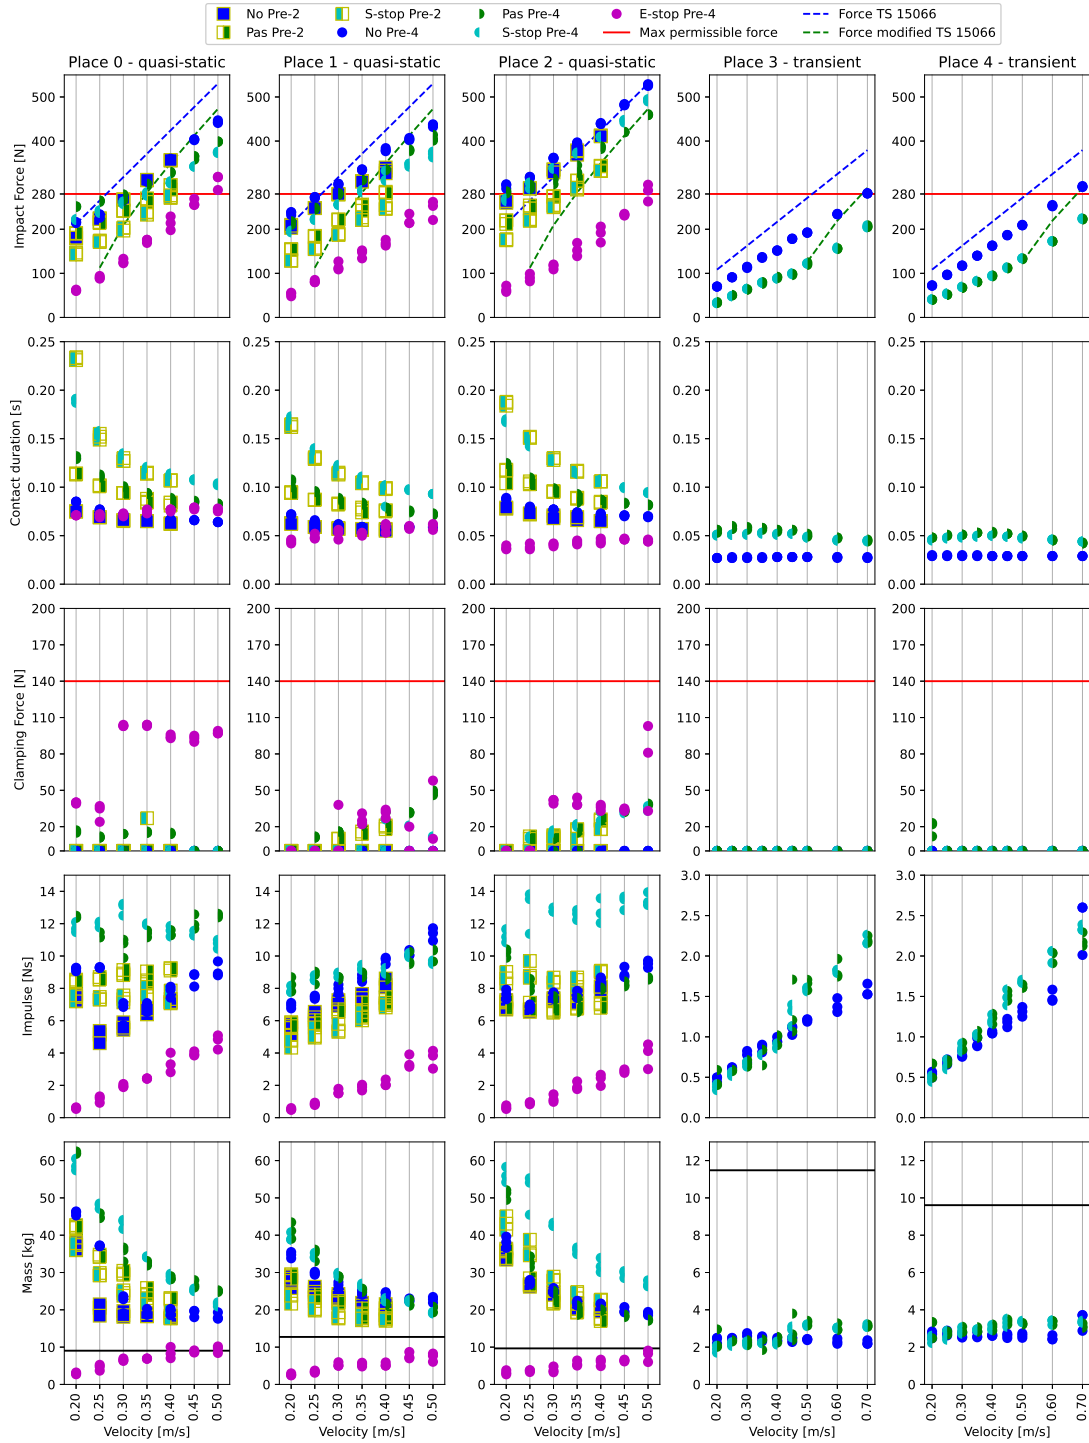


Figure 2.7: UR10e summary of measurements. The AIRSKIN module is either absent (‘No’), active (S-stop/E-stop) or merely pressurized but not active (‘Pas’). The robot is either in the least restrictive preset (‘Pre-4’) or in the second most restrictive preset (‘Pre-2’). In addition, active skin could trigger either the Safeguard stop (S-stop) or the Emergency stop (E-stop). The duration is the time between the impact detection and the first local minimum of the force measurement, i.e., the end of Phase I. The maximum permissible force is derived from ISO/TS 15066. Dashed lines are force predictions from Eq. 2.6 – see Sec. 1.4. The effective mass is calculated from the UR10e model. Originally published in [59].

the impulse could be increasing with velocity or constant, for places 1, 3, 4, and places 0, 2, respectively. Combined with the previously seen contact duration, we can conclude that the skin extends the contact duration and distributes the impact energy in time. However, the constant impulses in Place 0 and 1 show that the skin presence also increased the overall transmitted energy.

The estimation of the robot mass significantly exceeded the computed effective mass. This can be even five times for UR10e, the estimation for Place 0 and 2 for low velocities is 60 kg instead of the measured 12.5 kg. This can lead to a hypothesis that the effective mass itself is not enough to determine the collision force characteristics and that the robot dynamics and controller behavior should also be considered.

These findings are also numerically summarized in Tab. 2.5. Therefrom we can make additional observations. For quasi-static contact and the same safety preset, the difference between the passive and active use of the skin with the UR10e can be as small as only 4 % (compare the mean for ‘Pas Pre-4’ of -6 % with ‘S-stop Pre-4’ of -10 %). The use of Pre-2 leads to lower impact forces in general.

Setup	Quasi-static				Transient		
	P0	P1	P2	Mean	P3	P4	Mean
Pas Pre-4	1	-8	-11	-6	-41	-39	-40
No Pre-2	-4	-10	-8	-7	–	–	–
S-stop Pre-4	-7	-16	-7	-10	-41	-38	-39
Pas Pre-2	-12	-29	-23	-21	–	–	–
S-stop Pre-2	-26	-39	-30	-32	–	–	–
E-stop Pre-4	-48	-59	-62	-56	–	–	–

Table 2.5: Mean difference of peak impact forces for UR10e compared to skin ‘No Pre-4’. ‘Pre-2’ stands for the second most restrictive and ‘Pre-4’ stands for the least restrictive safety preset. The AIRSKIN can either be absent (‘No’), merely pressurized but not active (‘Pas’) or active (E-stop/S-stop). ‘E-stop’ stands for the scenarios where AIRSKIN activated the Emergency stop while ‘S-stop’ means that the Safeguard stop was activated.

The reaction times studied on the UR10e are presented in Fig. 2.8 and 2.9. If AIRSKIN is allowed to trigger any kind of stop behaviors, then these are initiated at the contact with the skin and before a contact with the robot itself. However, if we rely only on the robot’s sensors (i.e., AIRSKIN is passive or not installed), the stopping behavior is initialized only after these sensors detect the collision. This means that if AIRSKIN is installed but passive, it already exerts pressure.

The comparison between Fig. 2.8 and 2.9 also shows that the passive properties of the pad affect differently the reaction times based on the impact place. The active pad and UR-skin lead to the same reaction times in all the places except for Place 2.

KUKA iiwa

Similarly to the UR robot, Fig. 2.10 summarizes the measurements with respect to five perspectives. The use of the Pad (passive or active) leads to lower impact forces compared to the impacts without the Pad (skin No). These effects are visible in both types of collisions, the quasi-static and transient, and in general lead to lower impact forces than the permissible force. Similarly to the UR10e, the transient collisions result in the same impact forces for both the active and passive Pad shown by the overlap of the semicircles, suggesting that only the




Figure 2.8: Reaction times for the UR10e using Pre-4. The AIRSKIN can either be absent (‘No’), merely pressurized but not active (‘Pas’) or active (E-stop/S-stop). E-stop stands for the scenarios where the robot used the least restrictive preset (‘Pre-4’) and AIRSKIN activated the Emergency stop instead of the Safeguard stop. Originally published in [59].




Figure 2.9: Reaction times for the UR10e with Pad and UR-skin. Reaction times are shown on the vertical axis while collision velocities are on the horizontal axis. The UR10e is either used without any skin or it uses UR-skin or the Pad in a passive regime (‘UR-skin Pas’, ‘Pad Pas’) or the skin activates the Safeguard stop (‘UR-skin S-stop’, ‘Pad S-stop’). Originally published in [59].

passive properties affect the impact force. However, unlike the UR10e, quasi-static collisions lead to clamping forces (see third row in Fig. 2.10) exceeding the maximum permissible force of 140 N even with the use of the Pad. The dashed lines—force predictions using Eq. 2.6—are consistent in the sense that the modified model (green lines) predicts lower collision forces for the elastic cover.

The collision durations are independent of the velocity and are constant. In addition, collision durations for skin No and other cases are small. As the duration stayed constant but the impact forces increased with the velocity, the impulse is also increasing. Interestingly, while impact forces for skin No were higher than for the other cases, the impulse is the same or lower.

The KUKA iiwa’s mass measurements show the significance of the impact location very clearly. While Place 0 matches the robot’s predicted effective mass, in the other two places the resulting effective mass exceeds significantly the predicted value. This is visible especially in the case of skin No.

The KUKA iiwa data present two different observations for quasi-static and transient cases in Tab. 2.6, top row. Both cases show the great influence of the passive properties of the Pad (a decrease of impact force of 32 % or even 46 %). While in the transient collisions the activation of the skin did not show any effect, it can decrease the mean impact force by another 6 % in the quasi-static case. There is no significant difference between the two variants of Stop 1.

The effect of AIRSKIN is more prevalent if the external torques activate only Stop 1 – see Tab. 2.6 bottom row. We see a lower effect of merely passive skin (average 26 % as opposed to 32 % with ext. torque on Stop 0) but a much larger improvement if the Pad activates a Stop 0 (average 67 % as opposed to 40 % with ext. torque on Stop 0).

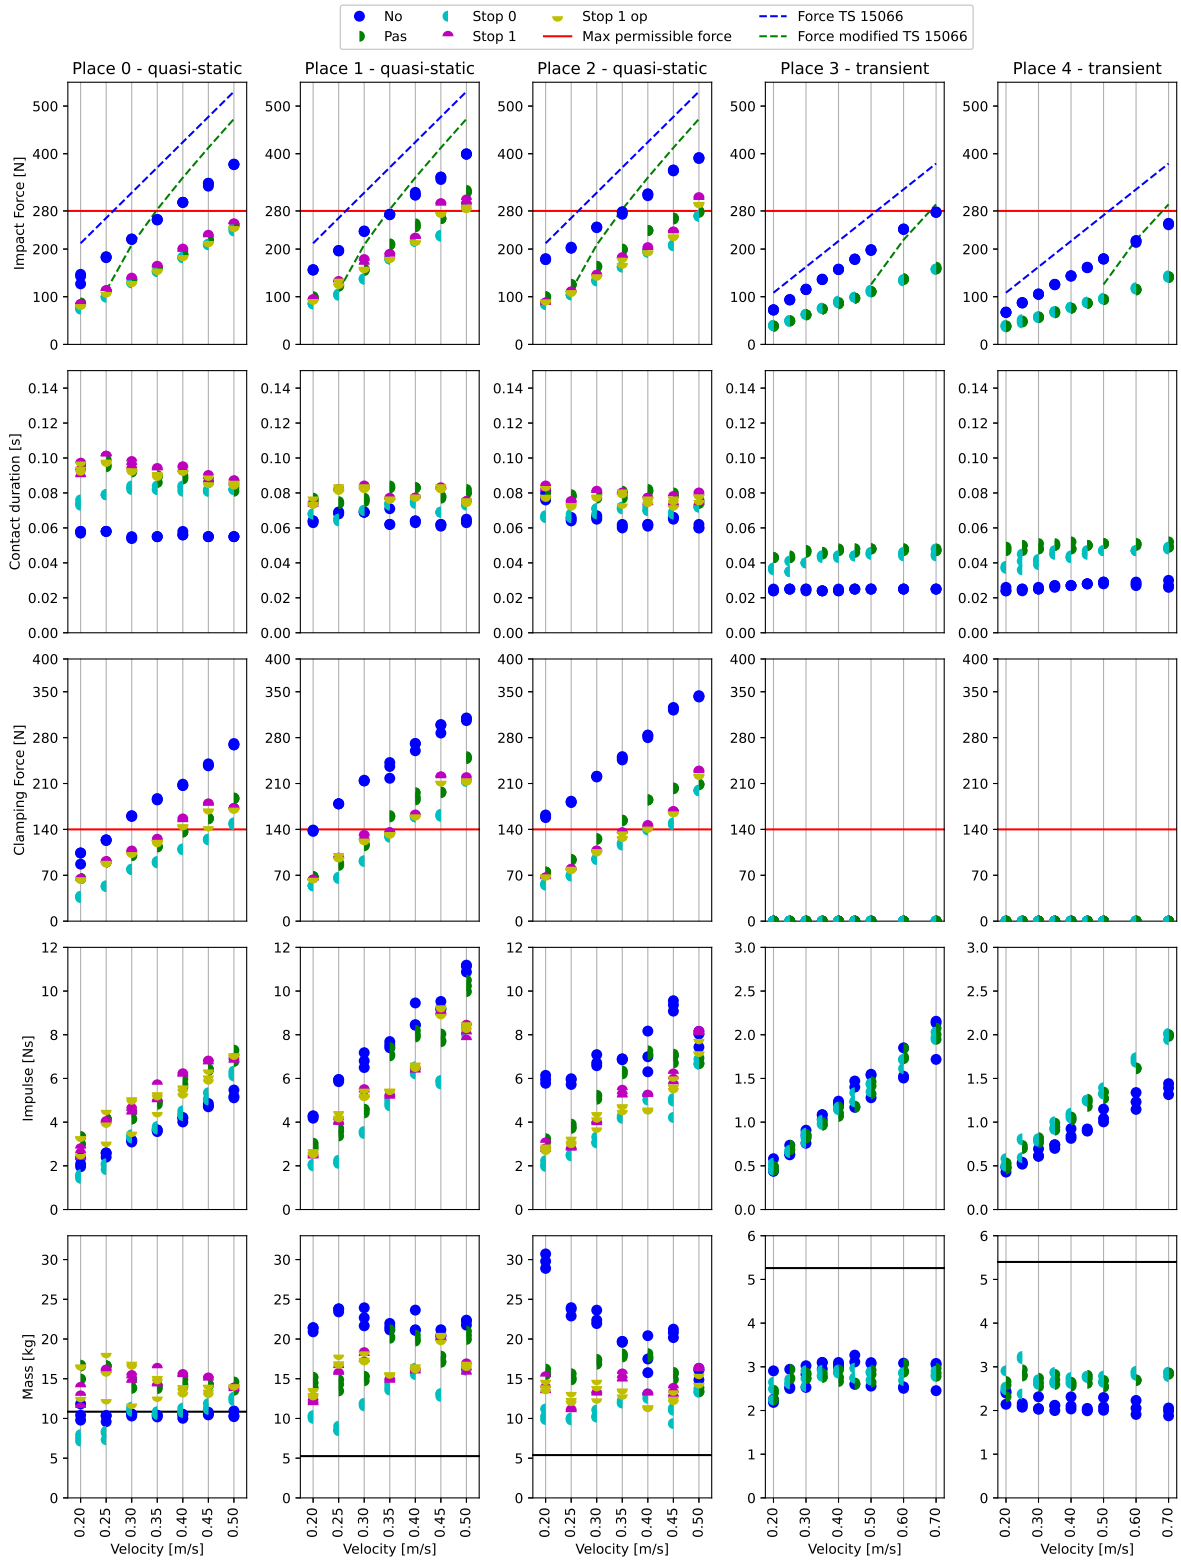


Figure 2.10: KUKA iiwa summary of measurements. The AIRSKIN module is either absent ('No'), merely pressurized but not active ('Pas') or active with various stopping behaviors ('Stop 0', 'Stop 1', 'Stop 1 op'). The maximum permissible force is derived from ISO/TS 15066. Dashed lines are force predictions from Eq. 2.6 – see Sec. 1.4. The duration is the time between the impact detection and the first local minimum of the force measurement, i.e., the end of Phase I. The effective mass is calculated from the KUKA iiwa model. Originally published in [59].

KUKA iiwa - Impact force change to skin No (%)

External torque limit on Stop 0							
Setup	Quasi-static				Transient		
	P0	P1	P2	Mean	P3	P4	Mean
Pas	-37	-28	-32	-32	-45	-46	-46
Stop 1	-36	-29	-38	-35	–	–	–
Stop 1 op	-38	-31	-38	-36	–	–	–
Stop 0	-41	-37	-43	-40	-45	-45	-45

External torque limit on Stop 1							
Pas	-41	-14	-23	-26	–	–	–
Stop 1	-48	-20	-33	-34	–	–	–
Stop 1 op	-47	-22	-35	-35	–	–	–
Stop 0	-66	-68	-67	-67	–	–	–

Table 2.6: Mean difference of peak impact forces for KUKA iiwa. The baseline impact force is the robot without any AIRSKIN modules. This value is compared with measurements where the skin is merely pressurized but not active (‘Pas’) or active with various stopping behaviors (‘Stop 0’, ‘Stop 1’, ‘Stop 1 op’). Originally published in [59].

KUKA Cybertech

Cybertech is an industrial robot with a significantly higher weight than collaborative robots and therefore results in different force profiles as shown in Fig. 2.11. The results present two diametrically different outcomes for quasi-static and transient collisions. Quasi-static collisions, the top row of Fig. 2.11, lead to clamping behavior without significant oscillation (denoted Type 3 in Sec. 1.3). Measurements could not be performed at Place 0 as the mounting of the pads on this robot did not permit to hit the measuring device with the downward movement. Velocities above 0.30 m/s lead to impact forces above 500 N at Place 1 too. Therefore we could not collect data to study trends as with the previous robots and focus only on the force profiles. Additionally, a special case is the collision at 0.40 m/s for Stop 2 where the impact moved the supporting table. The collected data show an interesting trend at Place 2. While at low velocities Stop 1 leads to lower forces, for velocities higher than 0.30 m/s, the pattern is switched, and Stop 2 leads to lower impact and clamping forces.

Transient collisions, the bottom row of Fig. 2.11, lead to very short contacts where even the robot without protective skin (‘No’) did not exceed the prescribed limit of 280 N. The use of the skin, nevertheless, significantly lowers the resulting impact force. A similar decrease as with the active skin can be achieved if Place 4 is used instead of Place 3. Also visible from the data is the prolonging effect of the soft protective cover when a collision with active skin can lead to a measurable force during a period almost twice as long as a collision without the skin (see Place 4).

Therefore, based on the transient collision data, Cybertech could be used in collaborative scenarios even without an active safety skin if it would be certain that only transient collisions without any clamping occur. However, any risk of quasi-static collisions, even for the low velocities, requires the usage of a device like AIRSKIN.




Figure 2.11: Time courses of measured force for KUKA Cybertech quasi-static (top) and transient (bottom) experiments for the $[0, 1, 0]$ and $[1, 0, 0]$ direction. No skin ('No') or with the active Pad initializing a specific stop category ('Stop 1 op' or 'Stop 2'). Originally published in [59].

Measured forces summary

Based on the results from particular robots presented above, we can summarize that the application of the skin leads to lower impact forces and the use of stricter stop categories leads to lower peak impact forces. In addition, we can observe an agreement across robots that in the transient collisions there is no difference in peak impact forces between active and passive skins (see Tab. 2.5, Tab. 2.6, Fig. 2.11). However, there are also differences between the robots as prominently visible on the different force profiles already presented in Fig. 2.2.

2.3 Stopping behavior effect on the impact forces

The effects of the stopping behavior settings for the collaborative robots are summarized in Fig. 2.12. The results are separated based on the robot-specific safety setting ('Pre-2' or 'Pre-4' for UR10e and using the external torque limit for the KUKA iiwa). The horizontal axis then captures the various stops triggered by AIRSKIN if it is present or the skin can be passive ('Pas') or absent ('No'). The measured peak impact forces are then represented with separate circles for each velocity increment (from 0.2 to 0.5 m/s with increment 0.05 m/s).

The vertical axis shows the various impact locations (Place 0–4).




Figure 2.12: Peak impact forces comparison for the various UR10e and KUKA iiwa with various external torque settings. The circles represent the measured 7 velocities (from 0.2 to 0.5 m/s with increment 0.05 m/s), where applicable. The initiated stop behavior is either a ‘Stop 0’, ‘Stop 1’, or ‘Stop 1’ op on KUKA iiwa. For UR10e, these stops were Safeguard stop (‘S-stop’) and Emergency stop (‘E-stop’) and also the safety preset was considered (‘Pre-2’ or ‘Pre-4’). The ‘Pas’ in the case of the AIRSKIN means the pads are pressurized but they do not initiate a stop, while ‘No’ means the AIRSKIN pad was removed from the robot. The three locations Place 0, 1, 2 are quasi-static collisions in the three directions downward, along y -axis, along x -axis respectively. The transient collisions along y -axis and along x -axis are Place 3 and 4 respectively. Originally published in [59].

The UR10e measurements in Fig. 2.12 and Tab. 2.5 also showed a significant effect of the specific stopping behavior of the robot. The skin improved the safety of the operation most when it was combined with the strictest stopping action or safety preset (‘Pre-2’ in our case). This effect was smaller with the KUKA iiwa (Tab. 2.6) if the external torque limit was active. However, without the external torque limit, the stopping behavior triggered by the skin became necessary to stop the robot.

Since we did not have the possibility to control all the stop categories for the UR10e, we investigated the effects of the various stopping behaviors only with the KUKA iiwa robot. All the KUKA iiwa quasi-static impact measurements are presented in Fig. 2.13 to demonstrate the effect of the various stopping behaviors on the final impact force in addition to Fig. 2.12. They are organized by the velocities into up to seven dots. In some cases, the exerted forces were higher than the measuring limit of our device (500 N) and thus we did not continue measuring for higher velocities or less safe setups (e.g., the ‘No’ skin setting columns). The KUKA iiwa robot allowed us to compare its external-torque-based stopping behavior with the AIRSKIN pad-based reaction between all three stop categories. The results support the earlier observation that AIRSKIN combined with a restrictive stop provides the best benefits.

Notice the first column where AIRSKIN triggers a Stop 0. It shows that the impact force stays low for all the measured velocities and even if the external torque trigger is not used. This finding was consistent across all the locations, see also Fig. 2.12. According to our data, AIRSKIN can serve as a replacement of external torque sensing. This is also the case for the Cybertech robot, see Fig. 2.11, where the more restrictive Stop 1 leads to lower impact forces.

Next to events generated by the active AIRSKIN, there may be safety events triggered by the collaborative robot itself. On the UR10e robot, collisions detected by the robot are handled internally by the robot controller. However, on the KUKA iiwa, the user can define how external torque limit violations are processed. Namely, one can choose whether this is connected to Stop 0, Stop 1, or Stop 1 op. This is visible in Fig. 2.13, where the peak impact forces for the three quasi-static impact places are presented with a combination of stopping behaviors. The settings of the skin are separate (horizontal axis). The skin would either be active and trigger the various stops, or be merely pressurized but passive ('Pas'), or it would not be present at all ('No'). The vertical axis captures the stops triggered by the KUKA external torque sensing capability or whether it was turned off ('Off'). The measured 7 velocities (from 0.2 to 0.5 m/s with increment 0.05 m/s) are shown as circles. Therefore, for example, the figure shows that if the skin triggers a Stop 0, the external torques have little effect on the impact forces (see the first column). However, this is not true for the inverse (see bottom row), as an external torque triggered Stop 0 can still lead to impact forces close to 400 N if the robot is not equipped with the Pad. However, if the Pad is used then the activity of the skin has little effect on the resulting forces.

The upper right corners in Fig. 2.13 do not contain measurements as with the given settings, the impact forces exceeded the 500 N limit of the measuring device. It is also visible that the impact places influence the resulting forces as at Place 1, the limit of 500 N was exceeded also with other stop combinations (see second column for Stop 1).

However, we can also notice the importance of the collision location from Fig. 2.13, namely the collisions at Place 1 (along the y -axis). While the general observations made so far still hold, we can notice that the exerted forces are larger with AIRSKIN. Similarly the impact places show an effect in Fig. 2.12 where Place 2, in general, registers higher peak impact forces.

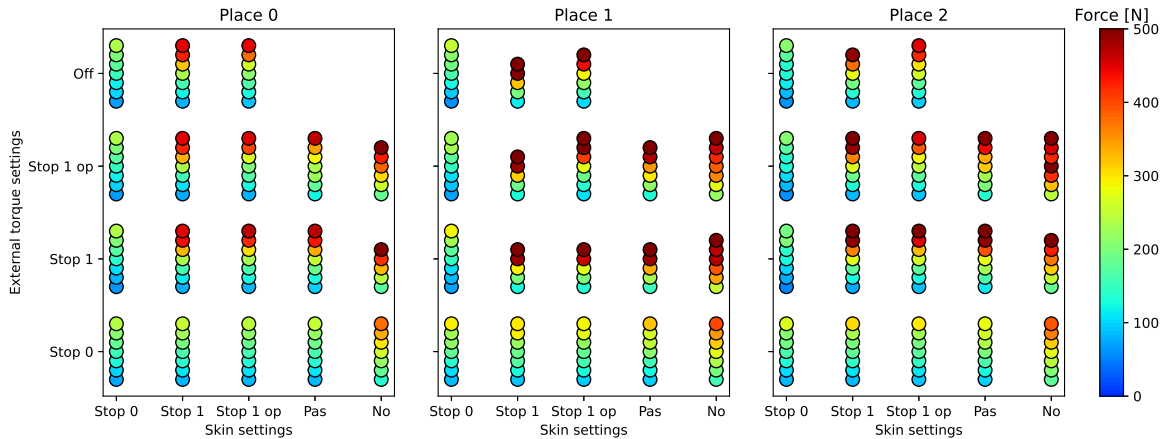


Figure 2.13: KUKA iiwa various stop combinations for quasi-static impacts. The initiated stop behavior is either a category 0 stop ('Stop 0'), a category 1 stop ('Stop 1'), a category 1 stop on path ('Stop 1 op'). The 'Pas' in the case of the AIRSKIN means the pads are pressurized but they do not initiate a stop, while 'No' means the AIRSKIN pad was removed from the robot. The 'Off' setting for the torques means that they were turned off. The circles represent the measured 7 velocities (from 0.2 to 0.5 m/s with increment 0.05 m/s), where applicable. The three locations Place 0, 1, 2 are downward, along y -axis, along x -axis respectively. Originally published in [59].




Figure 2.14: UR-skin threshold force (left) and stiffness (right). The measured values are color-coded. The impact point, B2, is also marked. Originally published in [59].




Figure 2.15: The Pad threshold force (left) and stiffness (right). The measured values are color-coded. The impact point, B2, is also marked. Originally published in [59].

2.4 AIRSKIN module properties

In this chapter, two types of AIRSKIN modules were employed: UR-skin (used on the UR10e robot, Fig. 2.4a) and Pad — AIRSKIN module pads employed on the KUKA robots (Fig. 2.4b). To assess to what extent these two versions affect the results, we studied two key properties: activation force threshold and mechanical stiffness. The point B2 marks the center of the impact (see Fig. 2.14 and Fig. 2.15).

The threshold force, i.e., the force at which AIRSKIN detects a collision, was 2 N for the UR-skin at point B2 and up to 8 N in the surrounding area (Fig. 2.14, left). The threshold force for the Pad was less than 2 N at point B2 and less than 4 N in the surrounding area (Fig. 2.15, left). Thus, in terms of threshold force, the two locations are comparable. In terms of stiffness, again, the measured values at location B2 are comparable (Fig. 2.14 and Fig. 2.15 right). Overall, UR-skin features much bigger differences in both properties across its surface, while the Pad is more uniform, with the exception of the central area that is stiffer (as there are electronics underneath). Impact force measurements with both types of AIRSKIN on UR10e (Tab. 2.7) show clearly lower average peak impact forces for the Pad compared to UR-skin, although the reaction times are longer (Fig. 2.9).

While the active use of both the skin and the pad outperform their passive use in the majority of quasi-static situations, this is not true for Place 2 where the opposite holds. The fast reaction time of the active AIRSKIN modules plays an important role in the overall performance as the UR10e starts to brake as soon a contact occurs with the skin (Fig. 2.8). This reaction time explains also the lack of difference between the transient impact forces for the passive and active variants. However, note that even a passive Pad performs better than the active UR-skin. Therefore the material used for the Pad plays a crucial role.

UR10e impact force change to skin No (%)								
Setup		Quasi-static				Transient		
		P0	P1	P2	Mean	P3	P4	Mean
UR-skin	Pas	4	-6	-10	-4	-41	-38	-39
	S-stop	-4	-16	-6	-9	-40	-38	-39
Pad	Pas	-5	-18	-26	-16	-45	-53	-49
	S-stop	-13	-27	-21	-20	-46	-54	-50

Table 2.7: Mean difference of peak impact forces for UR10e with the Pad and the UR-skin. The baseline impact force is the robot without any AIRSKIN pads. The values are then compared to either merely pressurized but not active (‘Pas’) or active skin triggering Safeguard stop (S-stop) modules. The robot always used the Pre-4 preset. Originally published in [59].

AIRSKIN pad type effect on the impact force

The measurements also support the importance of the skin’s passive properties, as visible in the comparison between the skin No and skin Pas columns in Fig. 2.12. The passive skin by itself can significantly lower the impact force, by up to 21 % (if paired with appropriate stopping behavior; see Tab. 2.5). This also supports the importance of the skin’s material properties of the used module. As visible in Tab. 2.7 and also Fig. 2.12, the Pad shows a significant improvement (even as large as the difference between having no AIRSKIN and using the URskin) in performance which can be traced back to its properties.

■ 3 Discussion and conclusion

We performed a total of 2250 measurements of impact forces in five scenarios and various velocities using two collaborative robots, UR10e and KUKA iiwa 7 R800, and the industrial robot KUKA Cybertech KR 20 R1810-2, all equipped with AIRSKIN safety covers. The dataset—involving collisions with active and passive protective skin as well as without it—is publicly available⁴. The main findings were the following.

First, we discuss transient collisions. Here, all the robots investigated showed similar behaviors. For the velocities used, even the non-collaborative KUKA Cybertech did not exceed impact collision force limits. We found significant effects of the passive properties of the protective covers—pressurized air in this case—on the transient collision. This was practically the only effect of the skin, as the active skin resulted in almost identical collision evolution and impact forces. Concretely, for transient contacts, passive safety covers resulted in 40 % lower peak impact force for the UR10e, 45 % lower for KUKA iiwa, and 46 % lower for KUKA Cybertech. However, with passive skin, the robot continues along its planned trajectory and clamping may eventually occur. With active skin, the robot will stop after contact is detected.

Second, we discuss quasi-static collisions. As expected, the peak impact forces are overall much higher. Unlike for transient collisions, the rest of the results has to be discussed taking the robot individual safety settings into account. In these cases, the effect of safety covers—active or passive—cannot be isolated from the collision detection and reaction by the robot itself. With passive safety covers, the contact duration is prolonged (200 % compared to no skin for UR10e; 25–50 % for the KUKA iiwa). On the UR10e, the effect on peak collision forces is moderate (decrease of 7 % for Pre-4 and 15 % for Pre-2). Still, the passive skin cover would allow to move from the more restrictive preset (‘Pre-2’) to a less restrictive (‘Pre-4’), while keeping the collision forces the same. On the KUKA iiwa, the situation is more complicated as we need to consider the connection of the external torque safety setting (see Tab. 2.6 and Fig. 2.13 and compare the columns ‘Pas’ and ‘No’). The biggest effects of passive skin padding are visible when the external torques are connected to Stop 0. In this case, the impact forces are lower by 32 % on average. For Stop 1, the effect is smaller (26 %). On the KUKA Cybertech, measurements with only passive skin were not possible.

With quasi-static collisions and active skin protection, the safety settings of the robot play a major part. For the UR10e manipulator, we found dramatic effects of the safety settings, namely to which safety stop category the skin was connected (Tab. 2.5). In the recommended setting, S-stop, the decrease in peak impact forces was 10 % for Pre-4 and 26 % in ‘Pre-2’ compared to no skin, but only 4 % and 13 %, respectively, compared to passive skin. However, the effect was much larger when connecting to the E-stop, namely 56 % compared to no skin and 53 % compared to passive skin. Note that this cannot be explained by the reaction times

⁴See at <https://osf.io/gwdbm>.

of the robot (Fig. 2.8), as these are similar for both stop types. It is the reaction of the robot that is responsible for the measured differences (cf. also Tab. 2.1). For the KUKA iiwa, one needs to consider the combination of safety settings: connections of ext. torque and the active skin (Fig. 2.13). First, Stop 0 clearly leads to the lowest impact forces. Second, the effects of external torque and skin safety settings are largely interchangeable. Moreover, they may even interfere with each other (e.g., ‘Stop 1 ext. torque’ with ‘Stop 1 op skin’ leads to higher impact forces than ‘Stop 1 ext. torque’ with ‘Pas skin’ or ‘ext. torque off’ and ‘Stop 1 op skin’). Third, there seems to be no measurable effect of the active skin compared to passive skin. Thus, for the KUKA iiwa, ext. torque safety settings seem to suffice to warrant safety. However, there may be applications where external torques occur naturally and hence active skin may still be needed for safety. On the contrary, on the industrial robot KUKA Cybertech, active skin is the only means to make the manipulator collaborative.

Next to empirical measurements, we also provided an extension of the simple collision model of ISO/TS 15066 that allows to consider the stiffness and compressible thickness of the protective cover (Sec. 1.4). Impact force predictions of this model are more in line with the measured data (Fig. 2.7 and 2.10 — first row) and this extension may thus be considered in future versions of collaborative robot standards for collisions with compliant surfaces.

Furthermore, we studied the force evolution after impact, schematically illustrated in Fig. 2.2. The first type, unconstrained dynamic impact, was observed in the UR10e robot. After collision is detected, the robot seems to actively retract, preventing clamping. The KUKA Cybertech displayed a constrained dynamic impact with clamping. Finally, the KUKA iiwa behaved similarly, but with an additional oscillation in the force profile. Importantly, the presence of passive or active skin protection did not alter this type of behavior (Fig. 2.6). The only case when the force evolution type changed on the same robot was on the UR10e when the skin was connected to the E-stop.

Let us compare our empirical results with what the collaborative standard ISO/TS 15066 prescribes (Sec. 1.3). For transient contacts, the maximum force limit is 280 N. In our experiments, this limit was never exceeded for all the velocities (from 0.2 to 0.7 m/s), not even for the non-collaborative KUKA Cybertech, where no collision was detected and no reaction triggered. Regarding quasi-static contacts, this is constituted by the force evolution during the first 0.5 s after impact—where the 280 N force limit applies—and the force evolution after the first 0.5 s, for which half of the force threshold, i.e., 140 N, applies.

Place	T [s]	TS 15066	mod. TS 15066	UR10e Pre-2			UR10e Pre-4			
				No	Pas	S-stop	No	Pas	S-stop	E-stop
0	<0.5	0.26	0.35	0.3	0.35	0.4	0.3	0.3	0.3	0.45
	>0.5	0.13	0.26	–	–	–	–	–	–	>0.5
1	<0.5	0.26	0.35	0.3	0.35	0.4	0.25	0.25	0.3	0.5
	>0.5	0.13	0.26	–	–	–	–	–	–	>0.5
2	<0.5	0.26	0.35	0.3	0.3	0.3	<0.2	<0.2	0.2	0.45
	>0.5	0.13	0.26	–	–	–	–	–	–	>0.5

Table 2.8: Maximum safe end effector velocities — UR10e [m/s]. Permissible velocities provided by ISO/TS 15066 and mod. ISO/TS 15066 taking skin compliance into account (see Eq. 2.5), and safe velocities determined from empirical measurements that do not exceed the collision force (280 N for $T < 0.5$ s, 140 N for $T > 0.5$ s) ‘No’ — no skin; ‘Pas’ — passive skin; ‘S-stop’ / ‘E-stop’ — active skin and its connection to robot safety inputs. Values in gray are minima of the corresponding rows for first and second phase after impact. Originally published in [59].

The corresponding maximum permissible velocities for the collaborative robots used here are calculated in Section 1.3 and presented in the “ISO/TS 15066” column of Tab. 2.8 (UR10e) and Tab. 2.9 (KUKA iiwa). The modified limits that take the stiffness and compressible thickness of the protective cover into account are in the “mod. ISO/TS 15066” column. Velocities that comply with the force limits computed from our empirical measurements are presented in the subsequent columns of these tables for the different settings. The maximum permissible velocity can be obtained as a minimum of the velocities for the first and second phase after impact.

Also note that for the UR10e robot (Tab. 2.8), with the exception of AIRSKIN connected to E-stop, no velocity limits result from the second phase (after 0.5 s). This is because despite the clamping nature of the scenario, no actual clamping occurs as the controller allows the robot to bounce back (see Fig. 2.6 and Fig. 2.2 — Type 1). Depending on the place in the workspace and the robot preset, robot velocities of 0.2 to 0.3 m/s can be safely operated. Addition of passive or active skin increases the safe velocity by approximately 0.05 to 0.1 m/s. If AIRSKIN is connected to the E-stop, velocities of 0.45 m/s become possible — contrasting with the ISO/TS 15066 prescription of 0.13 m/s (clamping scenario).

The situation is more complicated for the KUKA iiwa (Tab. 2.9) due to the optional involvement of the external torque limits and, importantly, because the second phase after the collision is present for this robot (see Fig. 2.6 and Fig. 2.2 — Type 2). When relying on the active skin only (ext. torque off), speeds higher than the norm prescribes (0.16 m/s for this robot and the clamping nature) can be safely operated. Namely, 0.2 to 0.25 m/s in the case of ‘Stop 1’ / ‘Stop 1 op’. A more significant productivity boost, 0.4 m/s safe velocity, constitutes a ‘Stop 0’ connection of the skin. Adding protective skin on top of ext. torque protection (right side of Tab. 2.9), brings about an increase in safe velocity from approx. 0.2 m/s to 0.3 m/s (passive skin) or 0.35 m/s (active skin).

Importantly, for both robots, the “mod. ISO/TS 15066” predictions (Eq. 2.5) are overall more accurate for the situations when protective skin is used than the original values from ISO/TS 15066.

Place	T [s]	mod.		KUKA iiwa external Torque limit off			KUKA iiwa external Torque limit Stop 0				
		TS 15066	TS 15066	Stop 0	Stop 1	Stop 1 op	No	Pas	Stop 0	Stop 1	Stop 1 op
0	<0.5	0.32	0.43	0.5	0.35	0.4	0.35	0.5	0.5	0.5	0.5
	>0.5	0.16	0.32	0.45	0.25	0.25	0.25	0.35	0.45	0.35	0.35
1	<0.5	0.32	0.43	0.5	0.25	0.3	0.35	0.45	0.45	0.4	0.45
	>0.5	0.16	0.32	0.4	0.2	0.2	<0.2	0.3	0.35	0.35	0.35
2	<0.5	0.32	0.43	0.5	0.35	0.35	0.35	0.5	0.45	0.45	0.45
	>0.5	0.16	0.32	0.45	0.25	0.25	<0.2	0.3	0.35	0.35	0.35

Table 2.9: Maximum safe end effector velocities — KUKA iiwa [m/s]. Permissible velocities provided by ISO/TS 15066 and mod. ISO/TS 15066 taking skin compliance into account (see Eq. 2.5), and safe velocities determined from empirical measurements that do not exceed the collision force (280 N for $T < 0.5$ s, 140 N for $T > 0.5$ s) ‘No’ — no skin; ‘Pas’ — passive skin; ‘Stop0’ / ‘Stop1’ / ‘Stop1 op’ — active skin and its connection to robot safety inputs. Values in gray are minima of the corresponding rows for first and second phase after impact. Originally published in [59].

In summary, our results demonstrate the following. For industrial robots (KUKA Cybertech) the effect of active protective skin is rather predictable. It allows for such a robot to detect collisions and respond, thus making human-robot collaboration possible. The passive

properties of the protective skin further decrease the impact forces. For collaborative robots, which have their own means of collision detection and response, the situation is more complicated and a number of other factors need to be considered. In this chapter, we studied the effect of robot velocity, safety settings, and position and impact direction in the workspace. While the effect of velocity is expected, the robot settings—where the skin and possibly other safety sensors are connected—were found to have important and sometimes intriguing effects. In addition, even the different AIRSKIN types, UR-skin vs. Pad, showed different effects on the impact forces. Thus, we conclude that empirical *in situ* measurements are needed and in particular, one should carefully consider the robot, its collision responses (force evolution), and the safety stops available. To illustrate this potential on the example of the UR10e robot, if collisions are the exception rather than the norm in an application and one connects the active skin to the robot E-stop, an almost four times higher velocity (0.5 m/s) will be safe, in contrast to 0.13 m/s that ISO/TS 15066 would prescribe.

Part II: Efficient Safe Collaboration

This part presents the research that combines two safety regimes, Power and Force Limiting (PFL) and Speed and Separation Monitoring (SSM), in order to improve the overall performance of the robot application. The first chapter presents the initial proof of concept that was enhanced upon in the later two papers, herein chapters. The common preliminaries are presented before the chapters themselves in the following section.

Performance research preliminaries

RGB-D cameras

Our implementation of the SSM regime relies on RGB-D cameras. The first chapter used an RealSense SR300 camera. The camera for the other chapters was Intel RealSense D435. We calibrate the robot and camera position through the ROS Hand-Eye calibration tool. The camera resolution is 848x480, and we use the RealSense short range presets⁵.

Human pose and velocity estimation

The computer collects two streams from the RealSense cameras: a color image aligned to the depth image (CAD) and a point cloud stream (PCS), also depth image aligned. The CAD image processing is done with OpenCV3 [120] running on a PC with a dedicated GPU. We use Intel RealSense SDK with PyRealSense to process these streams [119] to collect aligned color and depth images. The CAD image is sent the OpenPose library Python API [26, 121, 122] by PyOpenPose to estimate the human keypoints (see Fig. 3.2). Where either the COCO or BODY 25 model to identify the human keypoints. The resulting keypoint locations were then deprojected using the aligned depth image and thus we received the 3D coordinates of the operator's keypoints in the camera's frame of reference.

When human velocity was estimated, then the 3D velocity vector of the human keypoints was calculated from the 3D position change with simple differentiation and a moving average filter over a short past horizon to avoid jumps.

⁵See the file ShortRangePreset.json in the wiki pages at [119].

Chapter 3

Toward Unified PFL and SSM

This chapter is based on the publication:

Svarny, P.; Straka, Z.; Hoffmann, M.: Toward Safe Separation Distance Monitoring from RGB-D Sensors in Human-Robot Interaction In: *Proceedings of the International PhD Conference on Safe and Social Robots*. Strasbourg: Commission of the European Communities, 2018. p. 11-14.

In this work, we present a framework that realizes separation monitoring between a robot and a human operator in a detailed, yet versatile, transparent, and tunable fashion. The separation distance is assessed pair-wise for all keypoints on the robot and the human body and as such can be selectively modified to account for various interaction scenarios. The operation of this framework is illustrated on a Nao humanoid robot interacting in real-time with a human partner who is perceived by a RGB-D sensor.

Contributions

The contributions presented in this chapter are:

- We present a versatile and transparent separation monitoring approach.
- We demonstrate the tunability of this approach by experiments with human keypoint differentiation and dangerous tool differentiation.

The chapter is organized into three sections: Materials and methods (Sec. 1), Results (Sec. 2), and Discussion and conclusion (Sec. 3)

■ 1 Materials and methods

Human keypoints are perceived in the environment while robot keypoints are extracted from the model and current joint values. The relative distances are assessed and fed into the robot controller to generate appropriate responses.

1.1 Human keypoint 3D estimation

The collected keypoints from the process described in Sec. 3 are transformed into the Nao's frame of reference by affine transforms. The rotation and translation for them are gained from a pre-experiment calibration. We also used the model's confidence value to drop detections that were below 0.6 confidence as they were often false positives. This threshold was found by letting OpenPose analyze a scene without the human.




Figure 3.1: Experimental setup – collaborative workspace. (a) External view. (b) Camera view with human keypoint extraction. Originally published in [110].




Figure 3.2: Color image aligned with depth stream with the rendered human keypoints from OpenPose. Originally published in [109].

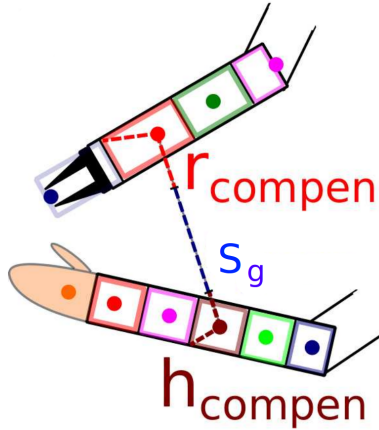


Figure 3.3: Separation distance calculation between robot and human keypoints. Originally published in [109].

1.2 Nao robot keypoints

A Nao humanoid robot (V3+) with keypoints on the left end-effector, forearm, and elbow was used to demonstrate the framework. We used forward kinematics with current joint encoder values as input to get the 3D position of these keypoints.

1.3 Separation distance representation

The *protective separation distance* S_p [11] needs to be maintained between any human and robot part such that the human will never collide with a moving machine. Its value will be determined based on reaction times etc. as in ISO/TS 15066. We extend S_p as a baseline with additional terms.

First, we want to account for “modulation” on the part of the human to grant larger distance from specific body parts (e.g. head) and on the part of the robot when carrying a sharp tool. Adding these distance offsets \mathbf{r}_s , \mathbf{h}_s gives rise to a *guaranteed minimal separation distance* S_g .

Second, as only distances between keypoints will be evaluated, but separation distance between any body parts needs to be maintained, we add compensation coefficients, $\mathbf{h}_{\text{compen}}$ and $\mathbf{r}_{\text{compen}}$ (see Section 1.4 below). This is the *keypoint separation distance* S_d —the quantity that will be monitored between any keypoint pairs.

Therefore S_d is in the form of a matrix of separation distances between two given keypoints i, j ($S_d^{i,j}$) (see Section 2).

$$\begin{aligned} S_g^{i,j} &= h_s^i + S_p + r_s^j \\ S_d^{i,j} &= h_{\text{compen}}^i + S_g^{i,j} + r_{\text{compen}}^j \end{aligned}$$

1.4 Keypoint compensation coefficients

Using a discrete distribution of keypoints allows fast calculation, but does not take the full volume of the bodies into account. The compensation coefficients r_{compen} and h_{compen} allow us to guarantee S_g even with a discrete keypoint distribution.

These coefficients are calculated in two steps. First, every part of the body is assigned to its nearest keypoint. Then the maximal distance over all of its assigned volume is selected as the compensation coefficient for the keypoint (see Fig. 3.3)—thereby always guaranteeing S_g .

1.5 Robot control

We used PyNaoqi to control the Nao. The Nao robot was moving his hands back and forth periodically in front of his chest. The robot stopped when an $S_d^{i,j}$ threshold was violated. The robot resumed operation upon “obstruction” removal. In addition, we defined a reduced speed distance: when $S_{d(\text{reduced})}^{i,j}$ was violated for any keypoint pair, the robot reduced its speed to half.

1.6 HRI setup

The Nao robot was sitting in a fixed position with respect to the camera that captured the robot’s workspace (see Fig. 3.2). Our setup is safe because of the Nao robot’s size and power. In a industrial setting with a potentially dangerous machine, S_p would be determined from ISO/TS 15066. In our case, the threshold was chosen arbitrarily.

End effector	Wrist	Elbow
0.06 m	0.05 m	0.06 m

Table 3.1: Keypoint distances for robot. Originally published in [109].

The compensation values accounting for keypoint density (Section 1.4) were determined by measuring the distances between keypoints (Tab. 3.1 and 3.2). Only upper body keypoints were taken into consideration for the human operator. We call the set of keypoints of the nose, neck, eyes, and ears as the human head. In both, human and robot cases, the compensation coefficients were symmetrical and thus we list keypoint pairs only once.

Nose	Neck	Eye	Ear	Shoulder
0.10 m	0.25 m	0.10 m	0.10 m	0.15 m
Elbow	Wrist	Hip	Knee	Ankle
0.15 m	0.15 m	0.00 m	0.00 m	0.00 m

Table 3.2: Human compensation values $\mathbf{h}_{\text{compen}}$. Originally published in [109].

2 Results

We performed three scenarios: (A) basic separation matrix, (B) specific separation values for the head of the human, (C) emulation of a sharp tool in the robot’s hand.¹ Distances between all human and robot keypoints were evaluated simultaneously online. However, for clarity, we present only the interaction of the robot end-effector with two human keypoints (the right wrist and the nose) in the plots below. The baseline protective separation distance was set to $S_p = 0.05$ m and the reduced speed regime $S_{p(\text{reduced})} = 0.20$ m.

2.1 Basic scenario

In the basic experiment, we monitored the distance between the human wrist and robot end-effector – see Fig. 3.4. The relevant separation matrices are in the Tab. 3.3.

Crossing the threshold into the warning regime is detected by the robot around $t = 7s$ as shown by the orange shaded area. The robot enters reduced speed mode at this point.

¹The video is available at <https://youtu.be/3DZyuuQ1qPo>.




Figure 3.4: Basic scenario: presented are Nao end-effector and human wrist keypoint distances and thresholds (S_d and $S_{d(\text{reduced})}$). Red area shows the stopping of the robot and orange shows the robot slowing down. Originally published in [109].

Robot \ Human	$S_{d(\text{reduced})}$	
	Nose	Wrist
End effector	0.36 m	0.41 m

Robot \ Human	S_d	
	Nose	Wrist
End effector	0.21 m	0.26 m

Table 3.3: Basic scenario: Separation matrix for keypoint pairs from Fig. 3.4. Originally published in [109].

Similarly, the next crossing is marked by red shading and the robot stops. The removal of the wrist from the safety zones resumes the robot’s operations.

2.2 Head and body discrimination

The \mathbf{h}_s for the head keypoints was enlarged by 0.15 m. This led to the robot’s higher sensitivity to situations when the human operator approached the robot with his head, as shown in Fig. 3.5.

In the first half of the experiment, we see the reaction of the robot to the wrist keypoint by slowing down. Later, we see that the robot reacts the same way to the nose keypoint, but at a greater distance than to the wrist. Notice the different reactions of the robot (shown by the different shading) for similar distances of the two keypoints. It merely slows down for the wrist but stops for the head keypoints.

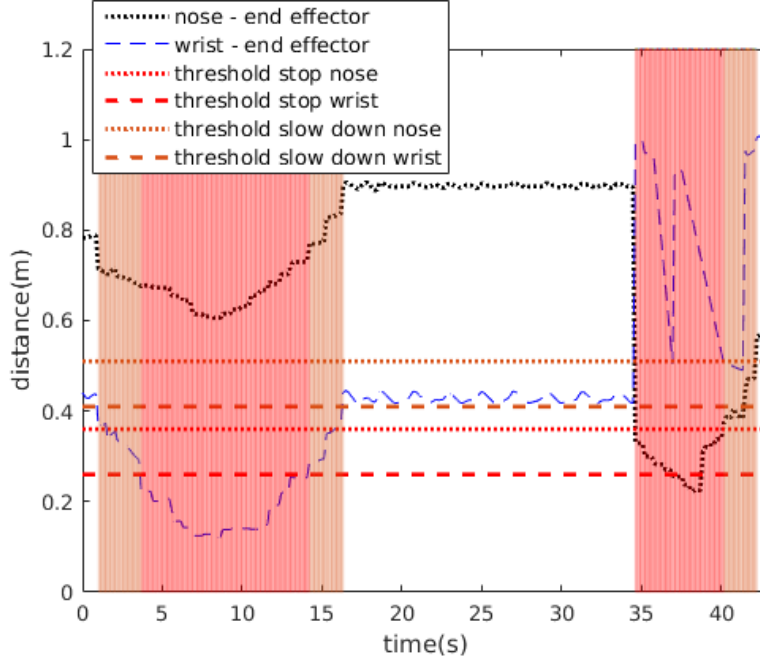


Figure 3.5: Head and body discrimination: A higher separation threshold for the human head region. Red area shows the stopping of the robot and orange shows the robot slowing down. Originally published in [109].

Robot \ Human	$S_{d(\text{reduced})}$	
	Nose	Wrist
End effector	0.51 m	0.41 m

Robot \ Human	S_d	
	Nose	Wrist
End effector	0.36 m	0.26 m

Table 3.4: Head and body discrimination: Separation matrix for keypoint pairs from Fig. 3.5. Emphasis is on values altered w.r.t. to the first scenario. Originally published in [109].

2.3 Dangerous tool usage

The left arm end-effector \mathbf{r}_s was increased by 0.1 m to simulate a possibly dangerous tool (see Fig. 3.6). The stopping and warning thresholds are now 0.1 m farther away from the robot end-effector. This increase is added to the original functionality from the head and body discrimination scenario, thus the robot reacts with greater sensitivity to the approach of the operator’s nose keypoint as opposed to the proximity of the operator’s wrist keypoint.

3 Discussion and conclusion

We presented a framework that realizes separation monitoring between a robot and a human operator. Distances are simply represented in Cartesian space in Euclidean norm and human and robot keypoints are treated equally. The separation distance is assessed pair-wise for all keypoints on the robot and human body and as such can be selectively modified. Velocity is not part of our representation but velocities can be converted into distance increments relying




Figure 3.6: Dangerous tool usage: Increased safety margin around robot end-effector. Red area shows the stopping of the robot and orange shows the robot slowing down. Originally published in [109].

Robot \ Human	$S_{d(\text{reduced})}$	
	Nose	Wrist
End effector	0.61 m	0.51 m

Robot \ Human	S_d	
	Nose	Wrist
End effector	0.46 m	0.36 m

Table 3.5: Dangerous tool usage: Separation matrix for keypoint pairs from Fig. 3.6. All values increased with respect to the previous scenario. Originally published in [109].

on measured quantities or worst-case constants per ISO/TS 15066. Also, we took velocity into account in our later research, see Chapter 5. The framework was illustrated on a Nao humanoid robot interacting with an operator monitored by an RGB-D sensor.

Chapter 4

Combining PFL and SSM

This chapter is based on the publication:

Svarny, P.; Tesar, M.; Behrens, J. K.; Hoffmann, M.: Safe Physical HRI: Toward a Unified Treatment of Speed and Separation Monitoring together with Power and Force Limiting In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019)*. 2019.

In this chapter, we extended the scenario from Chapter 3 and applied it on a KUKA LBR iiwa robot with the precise application of the ISO/TS 15066 safety regimes PFL and SSM in one collaborative pick and place scenario.

Contributions

The contributions presented in this chapter are:

- We deploy both regimes in a single scenario which provides in our view the unique contribution of this work at the time of publication.
- We use the vision pipeline capable of distinguishing different keypoints to adjust velocities based on the PFL regime prescribed thresholds for the given body parts.
- We present a cycle time-based performance study of various safety setups.

This chapter is composed of three sections: Materials and methods (Sec. 1), Results (Sec. 2), and Discussion and conclusion (Sec. 3).

■ 1 Materials and methods

1.1 Robot platform

A 7 DoF industrial manipulator KUKA LBR iiwa 7 R800 was used (see Robots in Chapter 5.4). The robot operates either at full speed (up to 1 m/s for the end-effector) or reduced speed (0.42 m/s). As an additional low-level safety layer, the KUKA *Collision detection* based on external torque estimation was turned on.

1.2 HRI setup

Our setup is illustrated in Fig. 3.1. A mock collaborative task has been staged: the robot performs a periodic operation. Operator periodically replaces one of the objects, entering the

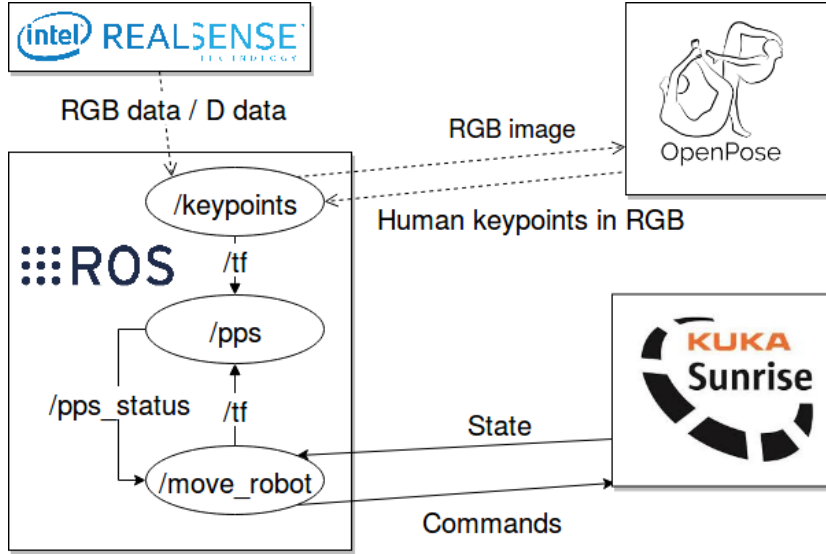


Figure 4.1: Software architecture schematics with the indication of ROS topics and transferred data. Originally published in [110].

robot workspace, and is perceived by the camera. The robot responds appropriately (slow down or stop). The robot was placed on a fixed table while the RGB-D sensor was on a fixed position so that it can capture the whole robot workspace. The camera was fixed to a construction that was separate from the robot’s platform to avoid tremors during the robot’s movement. The setup was designed to minimize the chance of occlusions.¹

1.3 Software framework and robot control

A schematics of the overall framework is shown in Fig. 4.1. OpenPose (see Sec. 1.4) finds human keypoints in pictures captured by the camera as orchestrated by a ROS node. The robot node consumes and produces information about the coordinate transformations. The relative distances are assessed in the peripersonal space module (*pps*) and fed into the robot controller to generate the appropriate response.

High-level control of the robot was done in the ROS node *move_robot*. We used the MoveIt! motion planning framework [123] to generate and execute the trajectories for our mock task. Our scenario additionally required speed modulation (stop, slow down, speed up) on the run which is not provided by MoveIt! and we have implemented a custom solution for smoothly modulating the trajectories in joint space, compliant with the corresponding limits of the platform. In brief, we used cascaded robot control which masks system non-linearities and lets us see the robot as a system of seven double-integrators, which we control similarly to a saturation controller [124]. We distinguish stopping motion and deceleration to reduced speed.

Stopping motion. The remaining trajectory of the robot is replaced by an alternative trajectory with a maximal deceleration for the fastest joint and relatively scaled deceleration for all other joints. The overall stopping time t_e is dependent on the velocity of the joints \dot{x}_j and the acceleration limits $a_{j,\min} \leq \ddot{x}_j \leq a_{j,\max}$, $t_{\text{stop},j}$ denotes the minimal stopping time for a joint j :

¹The complete setup including all experimental scenarios is illustrated in the accompanying video at <https://youtu.be/zP3c7Eq8yVk>.




Figure 4.2: Stopping motion using the trajectory controller. The stopping signal was received at 0.4077 s. The deceleration starts 0.0086 s later. The robot stops from the a speed of 1.1 rad /s in 0.3836 s (red area). The first red vertical line shows arrival of stop signal and the blue vertical line (0.0085 s later) marks the end of computation of the new trajectory. Note, that we consider the worst-case execution time in the selection of the reference state. Originally published in [110].

$$t_e = \max_{j \in Joints} t_{stop,j} \quad (4.1)$$

$$t_{stop,j} = \begin{cases} \dot{x}_j < 0 & \frac{0 - \dot{x}_{j,ref}}{a_{j,max}} \\ \dot{x}_j \geq 0 & \frac{0 - \dot{x}_{j,ref}}{a_{j,min}} \end{cases} \quad (4.2)$$

The worst-case run-time of the stopping trajectory calculation $t_{calc} \leq 0.02$ s was determined empirically. When the stop signal arrives, the earliest future state (with $t \geq t_{now} + t_{calc}$) along the current trajectory is selected and used as reference state \mathbf{x}_{ref} for calculations.

$$\dot{x}_j(0) = b_{j,1} = \dot{x}_{j,ref} \quad (4.3)$$

$$\dot{x}_j(t_e) = b_{j,1} + b_{j,2}t_e = 0 \longrightarrow b_{j,2} = \frac{\dot{x}_{j,ref}}{2t_e} \quad (4.4)$$

$$x_j(t) = \underbrace{x_{j,ref}}_{b_{j,0}} + \underbrace{\dot{x}_{j,ref} t}_{b_{j,1}} + \underbrace{\frac{\dot{x}_{j,ref}}{2t_e} t^2}_{b_{j,2}} \quad (4.5)$$

To facilitate the full breaking potential, we use polynomials (with parameters $b_{j,0}$, $b_{j,1}$ and $b_{j,2}$) of degree two to describe the joint positions. Hence, the velocities \dot{x}_j are linear with the maximum deceleration for at least one joint. This breaking behavior yields the shortest stopping time possible, but will for general trajectories slightly deviate from the original path. For point-to-point movements in free space (as in our example), this stopping strategy will remain on the planned path. Figure 4.2 shows the planned joint velocity and position, the stopping plan, and the joint velocity of a simulated robot.

Deceleration to reduced speed. When the signal arrives to slow down, a stopping trajectory is calculated as above. The original trajectory is scaled using the *IterativeParabolic-TimeParameterization* (MoveIt!) to comply with the desired reduced speed. When the linear

deceleration reaches the speed of the scaled trajectory, we search for the closest trajectory point ahead of the scaled trajectory. The scaled trajectory is shifted in time to continue after the deceleration and both trajectories are stitched at this point together. Acceleration back to full speed is performed similarly. The target joint position commands were then passed to the KUKA Sunrise cabinet via the FRI interface.

We took a conservative approach in the design of our controller as follows: when “pps status” signaled a more restrictive regime, it was executed immediately; conversely, in the other direction, a filter was applied to warrant that the operator has left the area. The pipeline described above is not safety-rated and the high-level robot control is capable of performing a Stop Category 2 only.

1.4 Used human keypoints

In addition to the mentioned setup in Sec. 3, the color images were sent to the OpenPose library Python API [121] to estimate human keypoints. For OpenPose, we used the COCO model and with the net resolution matching the input images. We also used the model’s confidence value to drop detections that were below 0.6 confidence as they were often false positives. This threshold was found by letting OpenPose analyze a scene without the human. We did not render the image in the final application, only for debugging or demonstration purposes.

These keypoints are represented as reference frames and added to the ROS transform library (called *tf*). The *tf* package stores the relationships between different coordinate frames in a tree structure, allowing for calculation of the position of the human keypoints w.r.t. the robot’s keypoints by using the relation between their frames.

Our experiment takes into account only upper body and hip keypoints detected by OpenPose’s posture model (see Fig. 4.3b), namely keypoints 0–7 and 14–17. These are the most relevant keypoints to our application and assume standard behavior of the operator. What we consider for our experiment as the human *head* are the keypoints of the nose (0), eyes (14, 15) and ears (16, 17).

1.5 Keypoint “bounding spheres”

Discrete keypoints allow a faster calculation of distances and unambiguous interpretability of the system’s expected behavior. Nevertheless, they do not take into account the full occupancy of the bodies, which could lead to the underestimation of the real separation distance. This problem is especially relevant with sparsely placed keypoints.

We need to guarantee S_p , the protective separation distance [11]. For this purpose, we introduce compensation coefficients for the robot r_{compen} and the human h_{compen} .

The calculation of the compensation coefficients with given keypoints is divided into two steps. In the first step, every part of the body is assigned to its nearest keypoint. Then, for every keypoint, the maximal distance over all its assigned part (from the first step) is selected as the compensation coefficient (see Fig. 4.3d)—thereby guaranteeing the separation distance S_p in all cases. With increasing density of the keypoints, the compensation coefficients get smaller.

In our case, the robot compensation values were determined from the model of the robot. For the human, the values were assigned empirically based on the distribution of OpenPose keypoints (Tab. 4.1). The human operator was interacting with the robot only with his upper body and the lower body was not taken into account. The resulting bounding spheres are in Fig. 4.3 and the values are in Tab. 4.1.




Figure 4.3: Keypoints and bounding spheres representation (aspect ratio kept). (a) Stopping and stopping after reduced speed distances. (b) OpenPose keypoint distribution [26] with bounding spheres on the keypoints of interest. (c) KUKA LBR iiwa keypoints (picture source: KUKA LBR iiwa brochure) with compensation bounding spheres. (d) Schematic 2D separation distance calculation between robot and human keypoints. The compensation coefficients are the distances between the keypoints and the farthest point of the body that belongs to the body part near the keypoint. Originally published in [110].

EE	7	6	5	4	3	2	1	Base
0.01	0.11	0.15	0.15	0.15	0.15	0.15	0.14	0.10
	Nose	Neck	Eye	Ear	Arm	Elbow	Wrist	
	0.10	0.25	0.10	0.10	0.15	0.15	0.15	

Table 4.1: Robot r_{compen} and human h_{compen} compensation values in meters. Originally published in [110].

1.6 Protective separation distance

We determined the terms of the protective separation distance equation, given in Eq. 1, as follows:

S_h $(t_r + t_s) \cdot v_h$, where v_h is the default human walking speed (1.6 m/s) [11], t_r is the time it took the robot to react to a issued stop status (0.1 s), and t_s the time it took the robot to stop its movement: 0.43 s, thus $1.6 \cdot (0.1 + 0.43) = 0.85$ m;

S_r $t_r \cdot v_{\text{max}} = 0.1 \cdot 1 = 0.1$ m;

S_s $t_s \cdot v_r = 0.43 \cdot 0.5 = 0.22$ m;

C the setup did not allow the operator to enter the workspace without being detected: 0 m;

Z_d see the h_{compen} values from Subsection 1.5: 0 m;

Z_r the KUKA iiwa's repeatability value: 0.0001 m.

The time t_s was determined based on measured calculation times (0.005 s) and the maximal deceleration of the robot which was set to 1.5 rad/s^2 .

Using these values, we can calculate the S_p as in Eq. 4.6

$$S_p(t_0) = 0.85 + 0.1 + 0.22 + 0.0001 = 1.17\text{m} \quad (4.6)$$

1.7 Power and force limiting

The SSM regime prescribes that the robot stops before contact occurs. In our approach, we also allow the robot to slow down so that it can operate in the PFL regime, see below. We assume the end-effector exerts pressure on a surface area of at least 1 cm^2 .

We can calculate the maximal relative speed of the system for a transient contact given the surface and the robot weight. For this, we use the formula A.6 from ISO/TS 15066. This equation also asks for some preliminary calculations, like for example μ , the reduced mass for the two body system of the robot and the human operator. We summarize the calculation here. In order to ascertain absolute safety, we assume the worst case scenario, i. e., an impact in the chest. The values for m_h , p_{max} and k are taken from the appropriate tables in ISO/TS 15066.

$$m_r = \frac{M}{2} + m_L = \frac{23.9}{2} + 0 \quad (4.7)$$

$$\mu = \left(\frac{1}{m_h} + \frac{1}{m_r} \right)^{-1} = \left(\frac{1}{40} + \frac{2}{23.9} \right)^{-1} \quad (4.8)$$

$$v_{\text{rel,max}} = \frac{p_{\text{max}} \cdot A}{\sqrt{\mu \cdot K}} = \frac{2.4 \times 10^6 \cdot 1 \times 10^{-4}}{\sqrt{\mu \cdot 2.5 \times 10^4}} = 0.50 \quad (4.9)$$

Thus we know that the speed of 0.42 m/s is a conservative speed in order to be in the PFL regime. We determine the distance at which the robot needs to start slowing down to

be PFL compliant in the same way as we did with SSM in Eq. 4.6. However, we take into account only the difference between 1 m/s and 0.42 m/s. The resulting value for S_p is 0.73 m (full to reduced speed). The stopping distance for 0.42 m/s according to the equation would be 0.60 m (reduced to stop). According to ISO/TS 15066, non-zero energy contact with the human head is not allowed. Thus our final setup forces the robot to stop on the proximity of the human head (see Section 2.3).

1.8 Keypoint separation distance representation

The separation distance is represented in a matrix of minimal effective separation distances for every pair of human-robot keypoints. This matrix allows the evaluation of the desired protective separation distance for all of the pairs individually. This matrix can be set explicitly or it can be a sum of different matrices as in our case.

The resulting separation distance is composed of several components—a *baseline* and any terms relevant from the safety perspective. The *baseline* is determined by the experimenter or calculated according to the methodology described together with Eq. 1 in Sec. 1.6. We have to evaluate the maximum possible speed and the protective separation distance based on the “worst cases over the entire course of the application” ISO/TS 15066. The resulting keypoints S_p^{ij} are added to compensation coefficients based on the bounding spheres $\mathbf{h}_{\text{compen}}$ and $\mathbf{r}_{\text{compen}}$ described already in Sec.1.5. This addition leads to the keypoint separation distances $S_{\text{kp}}^{i,j}$ between any two given keypoints i, j .

$$S_{\text{kp}}^{ij} = h_{\text{compen}}^i + S_p^{ij} + r_{\text{compen}}^j \quad (4.10)$$

Thus we calculate the keypoint separation distances for each keypoint pair. We show two calculations:

Pure SSM. According to SSM, the values necessary for a cat. 2 stop from full speed based on the Eq. 4.6 with the addition of the compensation values from Tab. 4.1 according to Eq. 4.10 are shown in Tab. 4.2 (left).

Combination of SSM and PFL regimes The robot first slows down and then stops only if needed. We add the calculations from Section 1.6; the resulting values are in Tab. 4.2 (middle). An example is provided in Eq. 4.11 with the nose-end-effector keypoint pair. Reduced speed is triggered at the distance $S_{\text{reduced,kp}}^{i,j}$ that is composed of $S_{\text{fulltoreduced}}$ per PFL (Section 1.7) and $S_{\text{reducedtostop,kp}}$ per SSM (Section 1.6, Tab. 4.2, last column).

$$\begin{aligned} S_{\text{reduced,kp}}^{i,j} &= S_{\text{fulltoreduced}} + S_{\text{reducedtostop,kp}}^{i,j} \\ 1.44 &= 0.73 + 0.71 \end{aligned} \quad (4.11)$$

Because of the shape of the KUKA robot, the values result in similar effective S_{kp} ; accordingly we list three keypoints from the robot and omit duplicate keypoint-pair values.

2 Results

The robot performs a mock pick-and-place task; the operator periodically replaces one of the objects, entering the robot workspace. The robot responds appropriately by slowing down or stopping and resumes operation whenever possible. The scenarios contrast the standard approach of a zone scanner or safety mat (Sc. 1, 2) with the pairwise distance evaluation between operator and robot keypoints (Sc. 3-5). Some scenarios employ a safe reduced speed

	Stop full	from speed	Reduce	speed	Stop reduced	from speed
	Nose	Wrist	Nose	Wrist	Nose	Wrist
End-effector	1.28	1.33	1.44	1.49	0.71	0.76
3	1.33	1.38	1.49	1.54	0.76	0.81
Base	1.28	1.33	1.44	1.49	0.71	0.76

Table 4.2: Effective keypoint-pair protective separation distance in meters. Originally published in [110].

per PFL (Sc. 2, 4, 5) and Sc. 5 issues a stop only on human head proximity. The description of the scenarios in our implementation (Sec. 2.1 – 2.3) is followed by a performance comparison on the mock task (Sec. 2.4). All upper body keypoints (see Fig. 4.3, right) were considered at all times, but we show only the safety-inducing keypoints in the plots below for clarity.

2.1 Scenario 1 and 2: Robot base vs. human keypoints

In the first two scenarios, the distances between the robot base and the human keypoints were considered. The baseline S_p of 1.17 m (Eq. 4.6) is extended by compensation coefficients specific to the human keypoint bounding spheres (Sec. 1.5, Tab. 4.2). In addition, as only the base of the manipulator is considered, the robot’s maximum reach of 0.8 m has to be added, giving 1.17+0.8 m, plus keypoint compensations.

In a similar manner, the second scenario approximated the setting with distance-based zones for reduced speed and stopping by using the values from Sec. 1.7. A reduced speed zone started at 2.13 m (0.73+0.6+0.8) and stop at 1.40 m (0.6+0.8). The separation distance for slowing down from the maximum velocity was a composition of the necessary distance for slowing down, the necessary distance to stop from the reduced speed, and the robot’s reach, see Fig. 4.3a.

2.2 Scenario 3 and 4: Robot vs. human keypoints

In Scenario 3, we measure keypoint-pair separation distance with respect to the robot’s moving parts (namely any joint above joint 3) to stop at $S_p = 1.17$ m. The fourth scenario involved a reduced speed zone (see Sec. 1.7). When a human keypoint got closer than 1.33 m to any of the moving robot keypoints, the robot slowed down. If the human got closer than 0.60 m, the robot stopped. The behavior of the system is illustrated in Fig. 4.4.

2.3 Scenario 5: Addition of keypoint discrimination

The last scenario described the case when the robot reacted with a stop only if the human head was closer than 0.60 m to the robot. Otherwise, the robot slows down (keypoint distance below 1.33 m). The behavior is illustrated in Fig. 4.5. Notice that the safety regimes of the robot were triggered by different keypoint pairs than in the case of the previous scenario in Fig. 4.4.

2.4 Performance in mock task

Here we quantitatively evaluate the performance on the task under the different safety regimes as described above. The robot performs the task 20 times (measured at one of the two target objects) and the time needed is recorded. As a baseline, we use the unobstructed task at




Figure 4.4: Scenario 4: Reduced speed (light area) or stop (dark) triggered by keypoint distances below threshold. Positions of selected joints showing the slowing down / stopping (continuous lines, right y-axis). Keypoint pair distances triggering the behavior are shown (individual data points, left y-axis). Relevant threshold values: Reduced speed at 1.63 m and the stopping behavior at 0.90 m. These values are based on Eq. 4.11 and the appropriate compensation values from Tab. 4.1. Originally published in [110].

full speed of the robot and at reduced speed. The full speed scenario would not comply with collaborative operation; reduced speed at all times would comply with the PFL regime, provided the operator head is protected.

The results are shown in Tab. 4.3. Operating the robot in the reduced speed PFL compliant regime, scenarios 4 and 5, outperformed most of the experimental scenarios. The scenarios that take pairwise distances between robot and operator keypoints into account and use two thresholds (scenario 4 and 5) performed better than all other collaborative regimes. The last scenario that stops only for the head keypoints achieves the best performance.

Full sp.	Reduced sp.	Sc. 1	Sc. 2	Sc. 3	Sc. 4	Sc. 5
154	256	267	254	257	231	228

Table 4.3: Task duration for different scenarios in seconds. Originally published in [110].

3 Discussion and conclusion

In this work, we used a robot in a mock collaborative scenario, in which it shares its workspace with a human. The operator’s position was perceived with an Intel RealSense RGB-D sensor and human keypoints were extracted using OpenPose. Our paper presents an application of the standard for collaborative robot operation ISO/TS 15066. The standard prescribes two




Figure 4.5: Scenario 5. See also caption of Fig. 4.4. As soon as the first threshold at 1.58 m is met, the robot reacts with slowing down. When the human operator crosses the second threshold at 0.85 m with his head, the robot stops. Thresholds contain the compensation from Sec. 1.5. Notice that the detection of the operator’s elbow below the threshold does not trigger a stop but it does lead to a longer reduced speed period. Originally published in [110].

collaborative regimes (SSM and PFL). However, to our knowledge there is no work considering both in a single application. We follow the standard to derive the protective separation distance (per SSM) and calculate the reduced robot velocity (in compliance with PFL constraints) and deploy them in a single framework. We demonstrate this union with an implementation of pairwise keypoint distance monitoring. Compared to classical zone monitoring, the keypoint distance method has higher resolution and constrains robot operation less. Also, keypoints can be treated differently, taking the sensitivity of human body parts or robot keypoints (e.g. sharpe edges) into account—in this way the constraints on collisions (per PFL) can be transformed into separation distances (per SSM).

The operation of this framework was illustrated with a KUKA LBR iiwa robot interacting with a human partner that is perceived by a RGB-D sensor during a mock collaborative task. Contrasting a classical “stop zone” from the robot base with the keypoint-based approaches confirmed the potential of the distance monitoring between pairs of keypoints.

Currently, occlusions could cause a misestimation of the human’s keypoint location and thus the distance. Possible compensations and thus future enhancements are to use multiple sensors, compensate for occlusion by creating a human model or filter out the robot body in the scene. With these additions, we could also incorporate active evasion of the human instead of our current reactive behavior (see [33]).

Chapter 5

Functional Mode Switching

This chapter is based on the publication:

Svarny, P., Hamad, M., Kurdas, A., Hoffmann, M., Abdolshah, S., Haddadin, S.: Functional Mode Switching for Safe and Efficient Human-Robot Interaction
In: *2022 IEEE-RAS International Conference on Humanoid Robots (Humanoids)*.
IEEE, 2022.

This paper presented the concept of robot functional mode switching to efficiently ensure human safety during collaborative tasks based on biomechanical pain and injury data and task information. Besides the robot's reflected inertial properties summarizing its impact dynamics, our concept also integrates safe and smooth velocity shaping that respects human partner motion, interaction type, and task knowledge. We further discuss different approaches to safely shape the robot velocity without sacrificing the overall task execution time and motion smoothness. The experimental results showed that our proposed approaches could decrease jerk level during functional mode switching and limit the impact of safety measures on productivity, especially when guided with additional task knowledge.

Contributions

Our core contribution is incorporating the intended interaction modes as input parameters of the robot control in addition to the usual inputs, e.g., human detection, distance, velocity. The integration of the intended interaction mode then allows us to modulate how safety is ascertained, i.e., SSM/PFL regimes and transition between them. We call this approach Functional Mode Switching (FMS). Our contributions can be summarized as follows:

- We suggest to use human-robot interaction modes (collaboration, cooperation, coexistence, autonomous behavior) as input parameters that modulate the behavior of the robot automaton and the use of industrial safety regimes (SSM and PFL).
- We introduce a smooth velocity shaping method that allows switching between various safety regimes and evaluation of this shaping under various criteria.
- We present a transfer of industrial safety experience to non-industrial settings.

This chapter is composed of four sections: Methodology (Sec. 1), Experiments (Sec. 2), Results and discussion (Sec. 3) and Conclusion (Sec. 4).

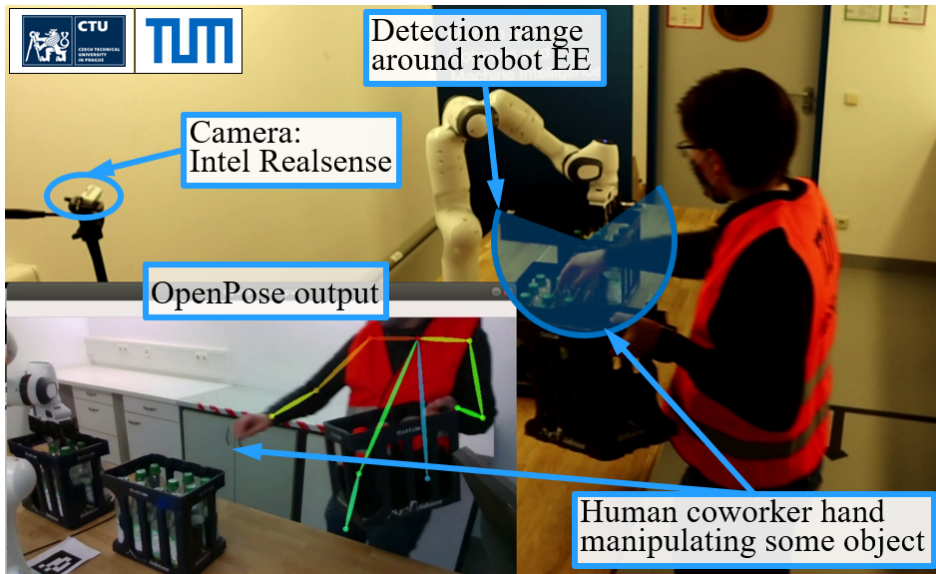


Figure 5.1: Experiment setup and OpenPose [122] view visualization. Originally published in [21].

1 Methodology

The setup of the application is presented in Fig. 5.1.

1.1 Relative velocity calculation

We use the Safe Motion Unit (SMU) framework [52] that calculates a biomechanically safe task velocity v^{SMU} based on the relative human-robot velocity and the appropriate force limits. Three collision-relevant situations can be distinguished based on the robot and human motions:

- (i) both partners are moving towards each other,
- (ii) both are moving in the same direction with the robot being behind the human and moving faster,
- (iii) the opposite case: the human follows the robot.

The effective speed of an impact v_{impact} is the highest in case (i). To mitigate the human injury risk, it must hold for the point of interest (POI), in our case end-effector (EE), velocity v_{EE} and human velocity v_H that¹:

$$v_{\text{impact}} = \|v_{\text{EE}} - v_H\| \leq v^{\text{SMU}}. \quad (5.1)$$

Regarding case (ii), the robot can be allowed to move with a velocity faster than v^{SMU} according to Eq. 5.1. However, if the human suddenly stops, the collision would be unsafe. Hence, v^{SMU} is used as the task safety threshold.

In case (iii), v_{EE} can be set equal to v^{task} , i. e., the velocity defined by the task. Velocity shaping cannot prevent collisions in this case as the robot would have to possibly exceed its desired value v^{task} to prevent a collision. Since the robot task velocity is commanded based

¹We use superscripts for commanded values and subscripts for observed or measured values, e. g., v^{SMU} is commanded. In contrast, v_{EE} is observed.




Figure 5.2: An exemplary interaction case in which the robot end-effector is moving in the direction \mathbf{u}_{EE} . Presented are the required position and velocity vectors for both the human ($\mathbf{x}_H, \mathbf{v}_H$) and robot ($\mathbf{x}_{EE}, \mathbf{v}_{EE}$) as also the projections onto the line connecting the human and end-effector locations, the desired task speed v^{task} , and safe velocity limit v^{SMU} . Originally published in [21].

on the relative velocity to the human, one may think it would be better to command a higher task velocity in case the human velocity exceeds v^{SMU} to escape collision. However, this would lead to a possibly unsafe robot velocity as in case (i). To assure a deterministic behavior of the robot in such cases, a practical solution is to set a lower limit for the task velocity resulting in a robot speed close to a standstill; we used $v^{\text{low}} = 0.1$ m/s.

We compared the following three shaping approaches with the baseline without human presence (v^{task} in our experimental validation).

- v^{motion} : relative velocity as projected velocity in the direction of the POI's motion,
- v^{short} : relative velocity as projected in the direction of the shortest distance between the human and the POI,
- v^{SMU} : uses only a distance threshold and shapes velocity based on the SMU's commanded velocity.

The difference between these shaping approaches lies in the way the relative velocity between the human and robot is being treated. Two approaches are given by the calculation of the relative projected velocity of the POI, i.e., point on the robot considered for the expected collision incident. In both these two cases, one does not use any projection. Moreover, the POIs can be chosen freely and could even encompass the whole body of the robot. The choice of a specific POI, however, does not change the principle of the presented approach as it can be recursively applied to all chosen POIs. Then the POI resulting in the most conservative robot velocity is to be adopted. For our study, we chose only the robot end-effector, EE, as a single POI.

A unit vector \mathbf{u}_c is determined based on the chosen projection, either for v^{motion} or v^{short} . We calculate the projection of \mathbf{v}_{EE} onto the chosen unit vector \mathbf{u}_c as:

$$\text{proj}_{\mathbf{u}_c} \mathbf{v}_{EE} = \left(\frac{\mathbf{v}_{EE} \cdot \mathbf{u}_c}{|\mathbf{u}_c|} \right) \frac{\mathbf{u}_c}{|\mathbf{u}_c|} \quad (5.2)$$

Therefore, for a chosen projection and thus velocity of choice v^c , the Eq. 5.1 must hold, i.e.:

$$v^c = \|\text{proj}_{\mathbf{u}_c} \mathbf{v}_{EE} - \text{proj}_{\mathbf{u}_c} \mathbf{v}_h\| \leq v^{\text{SMU}}. \quad (5.3)$$

1.2 Robot functional modes

In the following, we present our approach for safely and efficiently executing collaborative tasks. Given the biomechanical safety limits and task knowledge our approach avoids stopping the interaction (except for faults). Instead, the desired task velocity is only reduced as needed. In our work, we primarily use distance thresholds for switching between the functional modes. We use the relative distance d_{rel} that is measured between the human keypoint (e. g., the wrist $\mathbf{x}_{\text{wrist}}$) and a chosen robot POI (EE in our case), i. e., $d_{\text{rel}} = |(\|\mathbf{x}_{EE}\| - \|\mathbf{x}_H\|)|$, see Fig. 5.2. However, more complex switching behavior could be implemented (e. g., based on the task’s state).

Additionally, we introduce the *safe performance index SP*, i. e., the fraction of the velocity that should be used. The value of *SP* can be continuously updated online to reflect the desired safety level of the task execution. Its value can be derived from, e. g., percentage of braking distance of the robot, task and interaction knowledge user studies, etc. This index gives the human some control over the robot’s functional mode switching and task execution. For example, it can be set to force the robot to continue in the current mode or to switch into another mode by disregarding or overriding the automaton logical state transition functions based on task-relevant knowledge or partner experience. Adding this feature provides high flexibility for collaborative task execution under safety and performance considerations. Therefore, functional modes are switched based on the task specification or the current minimum distance between the observed points (unless there is a fault signal).

The collaborative task productivity may be lowered efficiently by switching to a safe robot operational mode only when necessary. By visually tracking the location of the human partner, flexible, near-real-time state transitions (below 1 ms) are ensured. For this, the robot dynamically switches between pre-specified functional modes that combine the interaction modes and safety regimes to increase the effectiveness of the interaction between the robot and the human partner. To provide a sufficient number of options for generic operational cases relevant to collaborative workspaces, we distinguish between four functional modes of the robot in pHRI scenarios (see Fig. 5.3):

1. Autonomous mode (*AM*),
2. Fault reaction mode (*FR*).
3. Coexistence mode (*Coex*),
4. Collaborative mode (*Col*),

Note that the *cooperative mode* from Introduction, Sec. 2.2 was omitted as it is not used during our example task.

Autonomous mode

In the autonomous mode, the task can be executed safely and autonomously while the human is outside the robot’s workspace. The robot carries out the task under specific optimality criteria, such as cycle time, leading to the full desired task speed to maximize productivity.




Figure 5.3: The robotic coworker automaton, based on [41]. Originally published in [21].

Fault reaction mode

The robot reaches a fault state. Therefore the robot motion is stopped until the fault is checked and cleared by the human partner and the task execution can be recovered again.

Coexistence mode

In this mode, the human and the robot share a common physical coexistence workspace, i. e., the relative distance between the robot and human is lower than a predefined coexistence threshold, $d_{\text{rel}} < d^{\text{coex}}$. Therefore, the desired task speed is lowered to ensure safety. The workspace is defined dynamically around the robot POI. However, no direct contact or interaction between the human and the robot is expected and therefore safety is provided by a combination of SSM and PFL. This is achieved by reducing the task velocity to respect the biomechanical safe limit v^{SMU} by optimal braking strategies, but only in cases where a collision is expected (see situations in Sec. 1.1).

Collaboration mode

Contrary to the no-interaction assumption for coexistence, collaborative tasks mostly involve closer physical interactions between the robot and human, i. e., $d_{\text{rel}} \leq d^{\text{col}} < d^{\text{coex}}$. The assumption of frequent human-robot contacts puts more focus on lowering the velocity, i. e., PFL regime, so that the relative velocity of a physical collision does not exceed v^{SMU} . The collaboration mode's triggering also depends on the current interaction scenario (see Sec. 1.1).

1.3 Smooth velocity shaping

We investigate velocity shaping, a necessity for coexistence and collaboration modes. Scenarios of direct pHRI, namely hand-over tasks as presented in [125], suggest that smooth, minimal-jerk movements on the robot side improve the task execution by the human partner. Therefore, any changes between robot functional modes and applied velocities should be smooth to make the interaction not only efficient but also pleasant. There are several ways to adapt the speed v with which a robot executes a preplanned motion profile while preserving both a safe and smooth trajectory. Based on task knowledge, key trajectory points of the underlying interpolator, and desired smoothness of the resulting motion, polynomial functions with higher




Figure 5.4: The basic idea of a general solution for the adaptive shaping of the robot speed using linear velocity blends. The dotted red lines indicate arbitrary possible target speeds v'_d , at which the shaping algorithm might leave the linear part. The values Δv_{start} and Δv_{stop} indicate the minimum change of speeds that can be carried out by the robot due to its acceleration. The times t_i, t^* indicate the onset of the acceleration and braking, respectively. Following these, the bold dashed line indicates the minimum increment in speed change given by the acceleration behavior. Originally published in [21].

degrees than the number of conditions to be satisfied can be used. For example, a fifth-order polynomial allows adapting the robot trajectory to arbitrary values with boundary conditions also in acceleration [126].

Assume that there are two given levels of speed between which the robot motion has to be adapted starting at time instant t_i (see also Fig. 5.4). The robot must transition from current speed v_i at position s_i to desired speed v_d at position s_d without violating the constraint on maximum acceleration a^m . To avoid jerky motions during velocity shaping, position, speed, and acceleration profiles along the direction of motion must be smooth. The shaping consists of three phases: raising the acceleration to its maximum value a^m , constant acceleration, and acceleration reduction to zero. Only the first phase is completely precalculated. The length of the other two phases is determined based on the current desired velocity v'_d – see dotted lines in Fig. 5.4.

■ 2 Experiments

For validation, we set up two main experiments. The “*Comparison* experiment” involves executing two tasks with different intended interaction modes. The goal is to compare the performance of various velocity-shaping approaches under different robot functional modes. The “*Grasping* experiment” demonstrates the safe execution of a generic collaborative grasping task using our proposed robot *Functional Modes Switching* concept with included task knowledge for enhanced efficiency. It is noteworthy that *Grasping* represents a common task in both industrial and non-industrial settings.

Note that all the described experiments were conducted employing the collaborative lightweight robot Franka Emika Panda and by the authors only.

2.1 Experimental setup

Two assumptions were used concerning the safe velocity limit v^{SMU} . First, calculating v^{SMU} is based only on available data from the head and chest collisions [52]. In [52], the authors also fitted a curve to the chest collision data and reached the following safety curve that we use too:

$$v^{\text{SMU}} = 0.1 \cdot (-0.4186 m_r + 5.2040), \quad (5.4)$$

with m_r being the instantaneous robot effective mass, lower cut-off at 0.1 m/s and upper one at 4.5 m/s. Second, as the experiment was performed on a robot already designed to be lightweight and safe, a scaling factor of 0.1 was therefore used for the safety curve to demonstrate the concept; see also [67]. Detailed motivation for the value of this scaling factor is beyond the scope of this thesis, but it is tied to the automotive industry origin of the collision data in the SMU framework [52].

We used the Franka Emika Panda Hand as a gripper, with a minimum contact area of 1 cm². Because the biomechanical data used in SMU are based on various shapes of POIs, the resulting v^{SMU} can be considered safe, especially after the application of the scaling factor.

The performance of the shaping approaches is compared in terms of relative productivity (calculated as the time needed to finish the task successfully) against the no-human case, i. e., executing the task with v^{task} .

2.2 Comparison experiment

This experiment compares pure coexistence and collaboration scenarios. We execute two tasks following the functional modes described in Sec. 1 to compare the performance of various velocity-shaping approaches under different robot functional modes. The robot end-effector was commanded to execute only a linear Cartesian movement to achieve higher peak velocities and easily interpretable data. Nonetheless, the methodology would also be applicable to more complex robot tasks.

The first task was a pure coexistence scenario between the human partner and the robot, i. e., where both human and robot share the same workspace but do not share the task purpose. However, the robot’s behavior is not affected by the human partner’s actions. The second task represented a real-time human-robot collaborative scenario with frequent activation of the robot’s *Col* mode.

We recorded two interactive human motions (one for each task type, *Data 1* and *Data 2*). For each recording, we ran the human motion cycles five times with given velocity shaping approaches as discussed in Sec. 1.1. The goal is to compare their resulting robot performance under the same human actions.

2.3 Grasping experiment

We consider a collaborative packaging scenario to demonstrate the usage of additional task knowledge to guide efficient safe execution. A bottle is picked up by the robot and moved to a target position while the human partner simultaneously grasps something in the robot’s work area, see Fig. 5.5. As a result, an unintended collision between the human and robot may occur, and appropriate mode switching is being performed. We ran the experiment four times with the prerecorded human tracking data. Similar to the first scenario, the robot used different velocity shaping approaches to complete the task.




Figure 5.5: Experiment progression in the work area with human-robot collaboration. Originally published in [21].

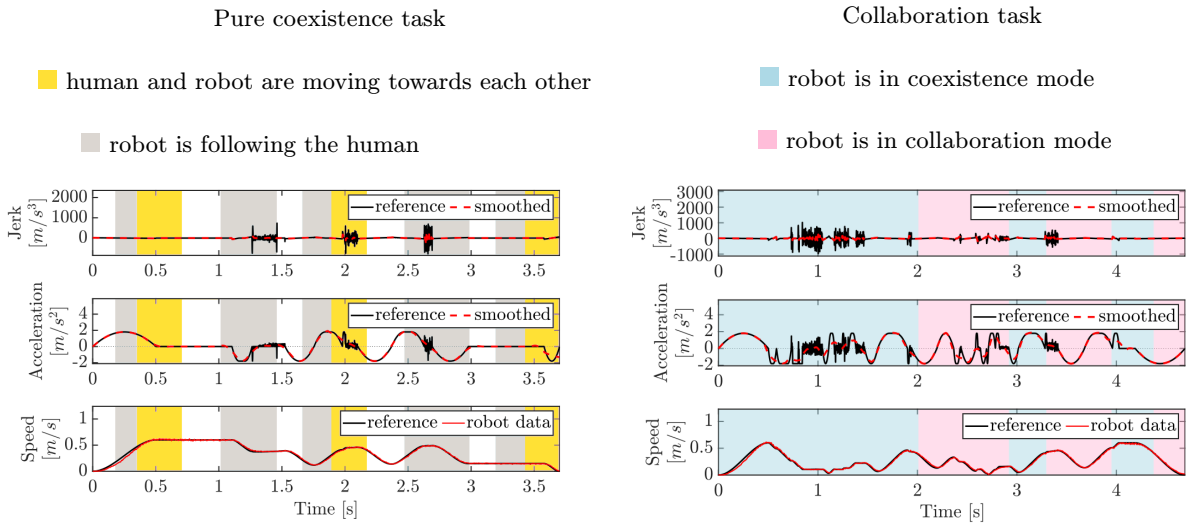


Figure 5.6: Safe task execution under two human-robot interactive scenarios: pure coexistence (left) and collaboration (right). The figures contain smoothed evolution of the noisy reference data (*smoothed*) and the actual measured *robot data* as opposed to the commanded *reference*. Originally published in [21].

Since part of the collaborative task involved vertical downwards robot motion (to grasp a bottle with possible human hand interference), we set the safety performance index as $SP = 0$ to enforce the safety-critical behavior for this part of the task. Otherwise, it is set as $SP = 1$. When no human was nearby, this grasping motion was executed with a full task speed of 0.5 m/s.

3 Results and discussion

3.1 Comparison of experimental results

Pure coexistence scenario

A typical robot speed motion profile using human position and velocity information for shaping the robot velocity in a pure coexistence scenario is shown in Fig. 5.6 (left). In this experiment, the relative speed between robot and human v_{rel} is shaped by changing the robot speed so that the *SMU* is only activated when needed and just as much as needed, respecting the biomechanical safety curve's limit. More specifically, for the shown profiles, the relative human-robot velocity is shaped when needed using the approach v^{short} . The calculated refer-

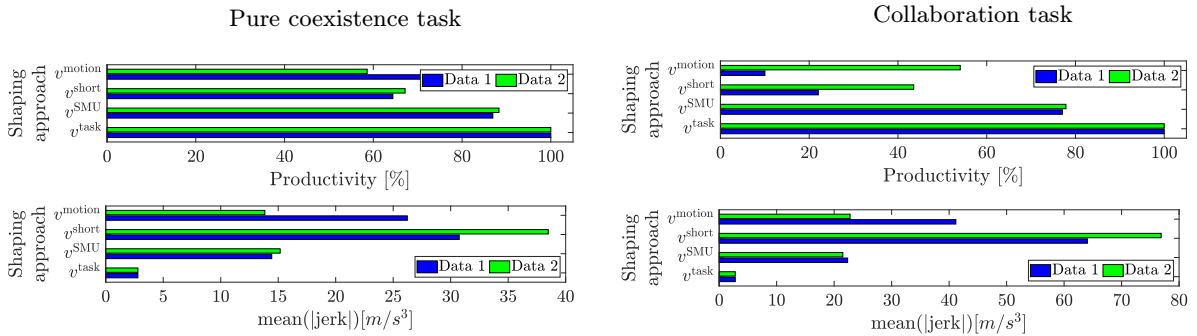


Figure 5.7: Task productivity (top) and average jerk levels (bottom) of different approaches for safe velocity shaping in the coexistence (left) and collaboration (right) scenarios as compared against the case when no human is around (i. e., the robot is in the autonomous mode). Originally published in [21].

ence acceleration and jerk profiles are also shown together with their smoothed averages. The robot’s velocity is continuously changing because the robot is responsive to human motion, and thus, velocity shaping is activated frequently. This makes the robot motion a bit jerky (e. g., at 2 s), but still far below the robot capability thresholds.

Both used approaches for relative velocity shaping during the experiments (i. e., v^{motion} and v^{short}) ensure human safety. When there is a collision risk, the human partner’s safety is guaranteed by the relative velocity respecting the used safety curve (Eq. 5.4). Both approaches result in nearly the same productivity, see Fig. 5.7 (top, left). However, Fig. 5.7 (bottom, left) shows that the v^{short} approach results in higher jerk on average. The stronger dependency on the human data for v^{short} is the probable cause of this jerk. Additionally, the approach v^{SMU} has a higher jerk on average than when using v^{task} , since the calculation of the v^{SMU} uses the configuration-dependent robot effective mass [63], so-called reflected mass $m_u(\mathbf{q})$ that changes during the task execution.

The human velocity data was noisy in the given setup. This could be a possible explanation for the lower performance of highly sensitive velocity shaping (v^{motion} and v^{short}) as opposed to the performance of the approach using only the distance information, i. e., the approach with v^{SMU} .

Collaborative scenario

A typical robot speed motion profile for the involved collaborative task, together with the recorded human interaction, is shown in Fig. 5.6 (right). The corresponding productivity and average jerk levels are also depicted in Fig. 5.7 (right). We use a simple *SMU* implementation in which only the human distance is used to trigger robot velocity shaping and this gives relatively shorter task execution times. As a result, the productivity is not that much sacrificed, it is at around 80% when compared to using v^{task} . However, this approach is oversimplistic and does not take into account the human speed v_H , which leads to unsafe situations if v_H is increasing in the robot motion direction \mathbf{u} as the resulting collision velocity would exceed the safe relative velocity.

Using the approach v^{motion} gives the worst performance in terms of productivity as compared to the approach v^{short} . Overall, higher productivity is achieved at the cost of higher jerk levels that are caused by continuous velocity shaping.

3.2 Grasping experiment results

A robot speed motion profile for the involved collaborative task while using velocity shaping along its motion direction, i. e., v^{motion} , is shown in Fig. 5.8 (bottom row). Since part of the task involved vertical downwards robot motion to grasp a bottle with possible human hand interference (shaded red area), we make use of our safety performance index and set it as $SP = 0$ to enforce the safety-critical behavior for this part of the task.

3.3 Summary

The following task execution times were obtained when comparing different velocity-shaping approaches for the same prerecorded human motion. The reference execution time when using full task speed v^{task} was 6.17 s. Using only the relative human-robot distance information for safe velocity shaping (i.e., approach v^{SMU}) the task execution takes 7.20 s, whereas using approach v^{short} it takes 7.34 s. Finally, with the depicted v^{motion} shaping approach 7.30 s is needed to finish the grasping task safely.

Another interesting aspect this motion profile reveals can be spotted by inspecting different force profiles in Fig. 5.8 (top row) outside the vertical motion zones. Using the suggested ISO/TS15066 impact model can result in forces higher than their specified tolerance limit. On the contrary, using a scaled safety curve with the proposed integrated safety scheme for human motion tracking and robot mode switching and control gives lower force levels.

■ possible human hand clamping ■ clamping-free vertical robot motion




Figure 5.8: Robot motion profile for the Grasping experiment with the inclusion of task knowledge of possible human hand clamping. A similar vertical robot motion without human interference is also shown for comparison. The figure shows the force limits given by the standard for the quasi-static or transient collisions. We differentiate cases where the human hand interference is excluded (shaded green area) or expected (shaded red area). Originally published in [21].



Figure 5.9: Example of a single arm application. Screenshot from GarMI concept video [20].

■ 4 Conclusion

In this work, we envisioned a situation where a humanoid robot would collaborate with the user on a task in close pHRI. It presented a transfer of experiences concerning safety of physical human-robot interaction from industrial settings to non-industrial robots, in addition to a method for an efficient combination of interaction modes and safety. This integration used the *Safe Motion Unit* scheme with human keypoint detection and a smooth velocity shaping controller. The proposed approach is evaluated with various relative velocity measures and expected interaction modes. However, it was found that higher task performance comes at the price of higher jerk levels for the robot velocity shaping while moving.

An industry-inspired collaboration scenario is demonstrated on a Franka Panda arm as shown in Fig. 5.1, where the human partner safety has to be always ensured. The results can, however, be extended also to a humanoid GARMi robot [19]. It is equipped with two Franka Panda robot arms and can perform similar tasks to the one we studied, see Fig. 5.9².

²See additional examples in the concept video [20].

Part III: Perception and Representation

Appropriate collision avoidance strategies cannot be selected without a representation of the space surrounding the robot. This space, so-called perirobot space, can be represented in various forms with regards to the needs of the application or robot platform. Multiple questions come forth in this context, namely the way the space will be structured, how will it change and what are its geometries.

Risk assessment is part of the safety evaluation of a robot application [10]. However, this risk assessment relies on experience and general guidelines, not a formal approach. This chapter aims to provide a modeling approach of the perirobot space (PeRS) that forces an explicit interpretation of sensor data and allows a formal evaluation of the overall sensor coverage.

A key component is the form of the representation of the robot and human body parts, or in general the representation of obstacles. Drawing on the results of the computer graphics community (see [127] for a survey), this often takes the form of some collision primitives. These can be simple shapes like spheres [46] or more complex meshes [99] and can differ for the robot and the human. Zanchettin et al. [37] represent robot links as segments and humans as a set of capsules. These shapes can also have a temporal aspect and represent so-called swept volumes, i.e. zones where the human or robot moved. Of course, for safety to be guaranteed, the whole body of both agents should be represented and considering only the robot end-effector or human hands or head does not suffice.

Because the applied idea of safety originates in industrial robotics, the safe zones are centered on the robot platform, as were the original laser barriers or gates. The novelty of these zones is, as opposed to the older fixed safety measures, their ability to dynamically change their shape based on the robot or human velocity [31, 37, 93, 99]. However, the application of biologically inspired approaches shifts the origins of safety zones around the person based on her peripersonal space [47, 100]. The robot surrounding space approach presented later in this work is inspired by both of these perspectives.

While the robot positions can be derived using forward kinematics, the human operator needs to be perceived by some sensor setup. Traditional sensor setups with low outcome resolution, like laser gates, are not suitable for many scenarios of pHRI. They have a large effective footprint of the robot or they do not allow close collaboration tasks (like handing objects) between the robot and the human without a significant drop in productivity. Therefore, either technologies that allow to distinguish the human body in higher detail or with greater precision on short distances are used. Examples of the first case are combinations of RGB-D sensors and convolutional neural networks for human keypoint/skeleton extraction from camera images [26, 27], or full 3D human body reconstruction [128]. These technologies together—albeit currently not safety-rated—make it possible to perceive the positions of individual body parts of any operators in the collaborative workspace in real time. The alternatives are, among others, wireless sensor networks that track operators who do not wear any devices [106] or proximity sensors distributed on the robot, usually part of electronic skins (e.g., Bosch APAS robot).

The applied technologies can influence the structure of the perirobot space too. Euclidean

space can be replaced by a depth space approach to account for the occlusions and specific geometries of the field of view of a RGB-D sensor [33, 46]. Also, a well established approach in robotics is the use of the configuration space (joint space) (see [33] for an overview).

Chapter 6

Perirobot Space

This chapter is based on the publication:

Rozlivek, J., Svarny, P., Hoffmann, M.: Perirobot space representation for HRI: measuring and designing collaborative workspace coverage by diverse sensors In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023 (submitted, available as arXiv:2303.02367).

Contribution

This chapter proposes to use a unified occupancy-based modeling approach for the space surrounding the robot, the *perirobot space* (PeRS). This modeling allows the comparison of various sensor setups and the evaluation of their workspace coverage. An illustration of the scenarios studied in this work is in Fig. 6.1. Based on the related research, the contributions presented in this chapter are:


- We introduce a unified robot surrounding space representation approach.
- We demonstrate the main benefits of this approach: the delimitation of occlusions, explicit definition of sensor data representation, and a clear description of the space monitored by the sensors.
- We present a novel use of established metrics to evaluate coverage of regions of interest.
- We apply the approach on an optimal (multi-) sensor coverage task on four different scenes.
- We make a publicly available code repository with the perirobot space implementation¹.

This chapter is composed of four sections: Proposed approach in Sec. 1, Results (Sec. 2), Discussion (Sec. 4) and Conclusion (Sec. 3).


■ 1 Proposed approach

Our interest in the space surrounding the robot is motivated by safe HRI. The inspiration comes from the so-called peripersonal space [100] which was also described as a “margin of safety” when avoidance behavior is concerned. However, perirobot space aims to be a more general approach that is not limited by the human-likeness of peripersonal space (e.g., sensors are not limited only to the robot’s body). We first discuss the region of interest for the


¹<https://github.com/ctu-vras/perirobot-space>




(a) Scene 1, human reaching towards the robot.



(b) Scene 2, human leaning against the table with the hand.



(c) Scene 3, human leaning against the table and obscuring the human working without floor contact.



(d) Dynamic scene, a human (left) walking around the table and observing the working human.

Figure 6.1: Studied scenes represented as occupancy models. Height is color-coded in the simulation output.

proposed approach. Then we present the modeling and representation of sensors and the notion of perirobot space itself. Finally, we present the metrics used for its evaluation and the modeling approach.

1.1 Region of interest

We study always a given region of interest, the volume where HRI can occur. In our approach, we study one robot-centered and one human-centered region. The robot-centered region is defined simply as the (semi-)sphere centered at the robot whose radius is given by the robot's end-effector's maximal reach, see Fig. 6.2a. This definition of the 'robot' space signifies that for safe interaction it is necessary to cover the robot's full reach. We designate as the 'human' space the bounding box enclosing the human that is next to the robot, see space h in Fig. 6.2b. By this representation, it is meant that this space needs to be efficiently monitored in order to ascertain safe interaction.

The surrounding space can be also defined with respect to the task, e.g., focusing on the human partner. Following definitions from [130, Chapter 3], for example, we could also distinguish either the robot's workspace (i.e., positions reachable by the end-effector) or the




Figure 6.2: Two regions of interest with a single sensor—a camera—and its perception space m . We distinguish four types of the space in the region of interest: occupied (red), free (green), unknown monitored free space (striped), unknown monitored occupied space (grey) and not monitored free space in the region of interest (dotted). Note that the occupied space is merely for illustration purposes this large. In practice, only the surfaces of the captured objects are registered.

robot’s envelope (i.e., the total volume of space occupied by the robot during these positions).

1.2 Sensor modeling

However, all of the mentioned definitions do not consider the sensors and their effect. The monitoring of the surrounding space together with the robot’s properties determine how the robot would be controlled. For example, a sensor that should trigger the robot’s stop at full speed in order to prevent a collision needs to be placed so that it gives the robot sufficient time to brake. Therefore, sensor information should also be part of the representation of the surrounding space of the robot. However, sensor monitoring also introduces a new challenge—occlusions and especially sensor data representation.

The data from the sensors monitoring the robot surrounding space are always represented in some manner. This representation is a deliberate choice and is not defined only by the collected raw sensor data itself. For example, a line laser sensor can be both a switch (i.e., something passed in front of the sensor) or it can serve as a so-called profiler (i.e., measure the profile of the object under the laser).

Additionally, the representation of the sensor reading is specific to the robot, the task, and the available sensors. The choice of representation can lead to different control decisions for the robot. For example, the collision classification [7] for a pressure sensor on the robot can have at least two representations. A collision can be interpreted as a stop signal or an impulse to move away.

Our approach is to model the representation of the sensor readings as information about space occupancy. While safety-related research focuses on detecting the human, we aim

to represent the space surrounding the robot as a whole as it is perceived by the sensors. Therefore our approach represents any relevant sensor data as occupied space, be it a human, robot or an object. We can distinguish three main ways of sensor representation:

Naive representation. Pure occupancy information provided from the sensor.

Volume-based. The sensor information is represented as occupancy of a predetermined volume (e.g., pressing a floor pad leads to the assumption that the whole space above the pad is occupied).

Feature-based. Sensor data are used to determine features (often human keypoints) and their locations. This representation can be further extended by determining a bounding box around the human, the feature neighborhoods (e.g., surrounding spheres), or volumes between features (e.g., cylinders connecting human keypoints) which are all considered as occupied space.




Figure 6.3: Visualization of occupancy based representations. The space monitored by the sensors is in green, while obstacles are in blue. Originally published in [129].

We analyze five sensors in this study. Our assumption is that thanks to the simple representation, any relevant sensor could be represented in a similar manner. Namely each sensor fits roughly into one of the three types:

Area. Sensors activated by events in a given area in the workspace.

Range. Sensors activated based on the range from the robot.

Zone. Sensors safeguarding a given zone but not necessarily monitoring the zone itself, e.g., a gate monitoring the entrance to the zone.

All of the analyzed sensors are defined by their 6D poses and additional defining parameters. The sensors and their defining parameters provided in the parentheses are the following:

RGB camera. Pinhole camera model, simulated as ray-casting (field of view, resolution), without the depth information provided.

RGB-D camera. Pinhole camera model, simulated as ray-casting (field of view, resolution), see Fig. 6.3a.

LIDAR sensor. Arc of rays (field of view, range), see Fig. 6.3b.

Pressure pads. Defined by their area and active if there are any contact points, i.e., occupied voxels right above the pad, designates all space above the pad as occupied (dimensions), see ‘P’ in Fig. 6.3c.

Robot proximity cover. Inflation of the robot model volume (inflation value), see ‘RP’ in Fig. 6.3c.

The combinations of sensors and representations used in this paper are listed in Tab. 6.1.

Sensor	Type	Representation		
		Naive	Volume	Feature
RGB Camera	Zone			x
RGB-D Camera	Range	x		x
LIDAR	Range	x		
Pressure pad	Area		x	
Proximity	Range	x		

Table 6.1: Representations and types of sensors. Originally published in [129].

1.3 Perirobot space

The *perirobot space* (PeRS) arises from this combination of a region of interest and occupancy-represented sensor data. A simplified representation of PeRS is in Fig. 6.2. This schema shows the monitored space m of the camera and two regions of interest, the robot and human space. Notice, however, that the robot itself changes what space is monitored by the camera due to occlusions, and thus the capability to represent its surroundings.

1.4 Coverage metrics

We chose two established metrics for classification tasks to evaluate the efficiency of the established PeRS in a specific region of interest, namely the F-score (F_1) and Cohen’s Kappa (κ). We distinguish three classes—positive (occupied space), negative (free space), and unknown (not monitored). The last class appears only in predictions. The F-score is given as:

$$F_1 = \frac{2TP}{2TP + FP + FN + UF + UO} \quad (6.1)$$

where UF and UO represent the not monitored free and occupied space, respectively. TP, FP, FN are defined in Fig. 6.4. The Cohen’s Kappa is given by the ratio of the relative observed agreement p_o and p_e , the hypothetical probability of chance agreement which is defined in multi-class version as:

$$\kappa = \frac{s(TP + TN) - \sum_k p_k t_k}{s^2 - \sum_k p_k t_k} \quad (6.2)$$

where s is the total number of voxels, p_k is the number of voxels where class k is predicted and t_k is the true number of voxels of class k .

1.5 OctoMap modeling

The core notion of our approach is occupancy in the surrounding space. All sensors are represented as providing information about the occupancy of the space they perceive. Specifically, we model the scene and, thereafter, all sensors using OctoMap [131] and octrees.

OctoMap and the occupancy representation also allow us to model the limitations of various approaches. For example, camera occlusions are considered because ray-casting from the simulated sensor is blocked by occupied space. Ray-casting also simulates the limited field of view of the cameras. Additionally, we introduced noise for the range-based sensors (LIDAR, RGB-D camera, and proximity cover) added to the measured distance.

As mentioned earlier, feature-based sensors create a representation of the detected human keypoints. We distinguish three models of the detected human keypoints based on the usual practice in safe HRI. Namely, a bounding box encompassing all the keypoints, the detected keypoints enclosed in spheres or cylindrical connections between the keypoints. These are shown in Fig. 6.5.




Figure 6.4: Example of the coverage metric representation for a monitored volume with two pads. The active pad designates all the volume above it as occupied. Only the voxels occupied by the person are truly occupied (TP, white on red), the voxels not occupied by the person are falsely marked as occupied (FP, red). The volume above an inactive pad is considered empty. The majority of the voxels are truly empty (TN, green), but the human arm voxels are falsely considered as empty space (FN, white on green). Last, there is also space not monitored by the pads containing empty space (dotted) and occupied space (white on dots). Originally published in [129].

Our modeling approach consists of the following steps:

1. Create a model of the scene and determine the region of interest.
2. Model sensors in the scene.
3. Calculate the sensor-generated occupancy.
4. Calculate the metrics. (see Sec. 1.4)

In our initial approach, we model a single manipulator placed on a table and a human in various poses next to it. As mentioned earlier, we investigate two regions of interest. While these could be determined in a task-dependent fashion, we chose regions that can be defined in a straightforward way for demonstration purposes.

■ 2 Results

To demonstrate our approach, we present three experiments (parameters in Tab. 6.2 evaluated in a series of HRI scenes (shown in Fig. 6.1). Although the presented scenes focus on HRI, the problem can be understood in general, as the robot’s workspace can contain unforeseen objects that could harm the robot or be damaged by the robot.

2.1 Different data interpretation experiment (Exp 1)

In this experiment, we compare the performance of the RGB-D camera observing the workspace based on different interpretations of the measured data. Those interpretations are: 1) 2D keypoint detection only—if a human is detected, the defined space around the robot is marked as occupied; 2) Raw 3D point cloud; 3) Raw 3D point cloud with added points from the keypoints surroundings (spheres); 4) Raw 3D point cloud with added points from keypoints cylinder connections; 5) Raw 3D point cloud with added points from the keypoints bounding box. The keypoint representations are shown in Fig. 6.5.

We placed the RGB-D camera in 172 different positions covering all four walls and ceiling (see Fig. 6.8 for all positions) and rotated the camera to five different orientations in each position, resulting in 860 measurement poses. We computed the F_1 and κ scores for each measurement pose for the ‘robot’ and ‘human’ spaces and evaluated the distribution of the metrics for each scene to evaluate the performance of each interpretation.

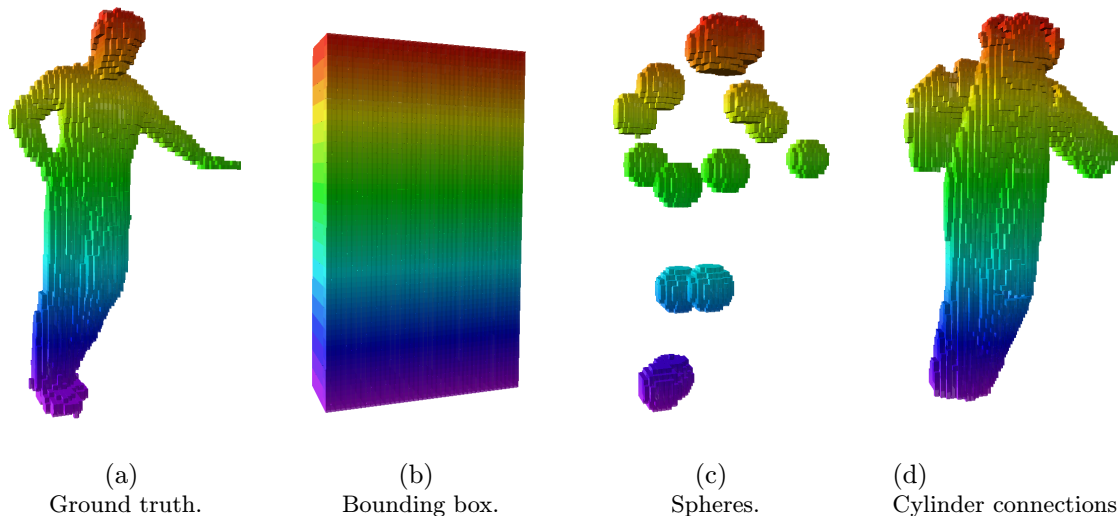


Figure 6.5: Ground truth and human models for the used keypoint representations. Height is color-coded in the simulation output. Originally published in [129].

As shown in Fig. 6.6 and Fig. 6.7, the trends are very similar for both metrics. The 2D detection is outperformed by others in both metrics and in both spaces. The results of the other interpretations indicate that the enhancement of the point cloud by keypoints improves the performance. The addition of cylinders has the highest median for both metrics and for both spaces. Moreover, for the ‘human’ space, the addition of cylinders has the highest maxima as well. For the ‘robot’ space, the addition of bounding boxes has the highest maxima. The addition of spheres has the worst results from the additions, but still outperforms the raw 3D point cloud.

In addition to interpretation evaluation, we can analyze the RGB-D camera placement for ‘robot’ space coverage. We created a heatmap, where each camera position is represented by its highest F_1 or κ value of the five orientations. Figure 6.8 shows the heat maps for the point cloud with the added keypoint cylinders in Scene 3. As can be seen, there are many solid placements based on the F_1 metric all around the walls. Moreover, it is clearly visible that placements close to the human that is obscuring the operator have a much lower F_1 score than the other placements around. The κ metric results emphasize the positions only on the right side of the perimeter with the best positions on the front wall. The simulated point clouds for the RGB-D camera placement with the highest ‘robot’ space values of F_1 and κ for Scene 1 and Scene 2 are shown in Fig. 6.9. The camera is placed on the right wall (in the same notation as for the heatmap) for the highest value of F_1 in Scene 1. In other cases, the camera is placed on the front wall, and the placement for the highest value of κ is the same for both scenes.

2.2 Multi-sensor coverage experiment (Exp 2)

The second experiment aims to evaluate the integration of several different sensors together and to determine how the combinations improve the coverage of the space. We integrated a pressure pad, a robot proximity cover, and an RGB-D camera placed on the robot base looking front. We tested 24 pressure pad positions, three proximity cover ranges (0.1, 0.2, 0.3 m), and 9 different orientations for the RGB-D camera (from -40° to 40° around the z -axis with a step of 10°). We evaluated the 7 possible combinations of the sensors—a triplet of sensors, 3 pairs of sensors, and 3 individual sensors. Unlike in the previous experiment, we

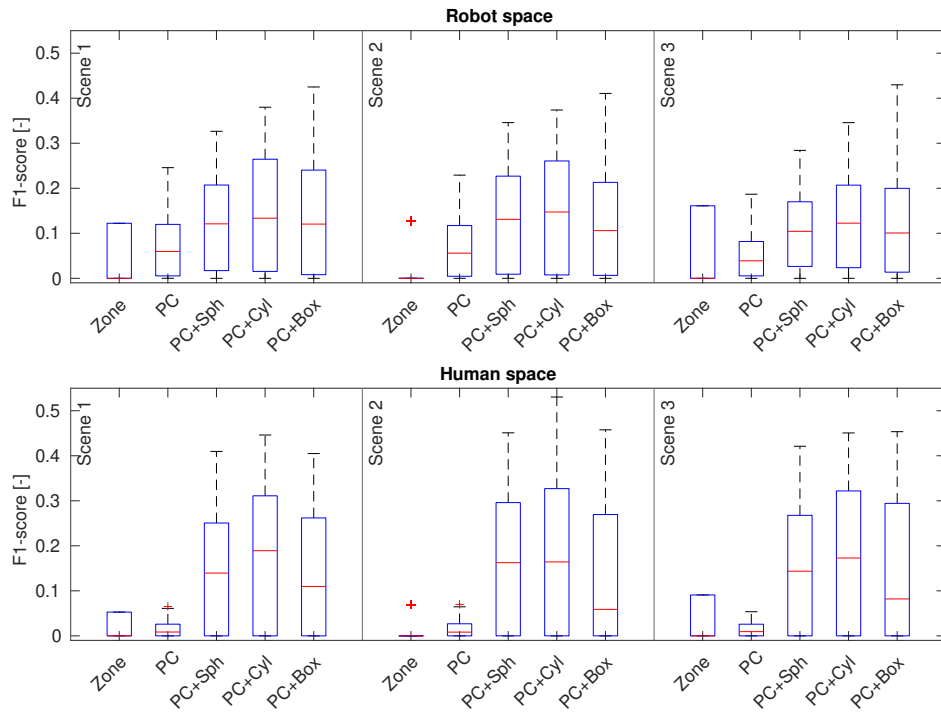


Figure 6.6: F_1 results in three static scenes for different interpretations of data from RGB-D camera—2D keypoint detection (Zone); raw 3D point cloud (PC); raw 3D point cloud with added spheres (PC+Sph), Raw 3D point cloud with added cylinders (PC+Cyl); raw 3D point cloud with added bounding box (PC+Box). Originally published in [129].

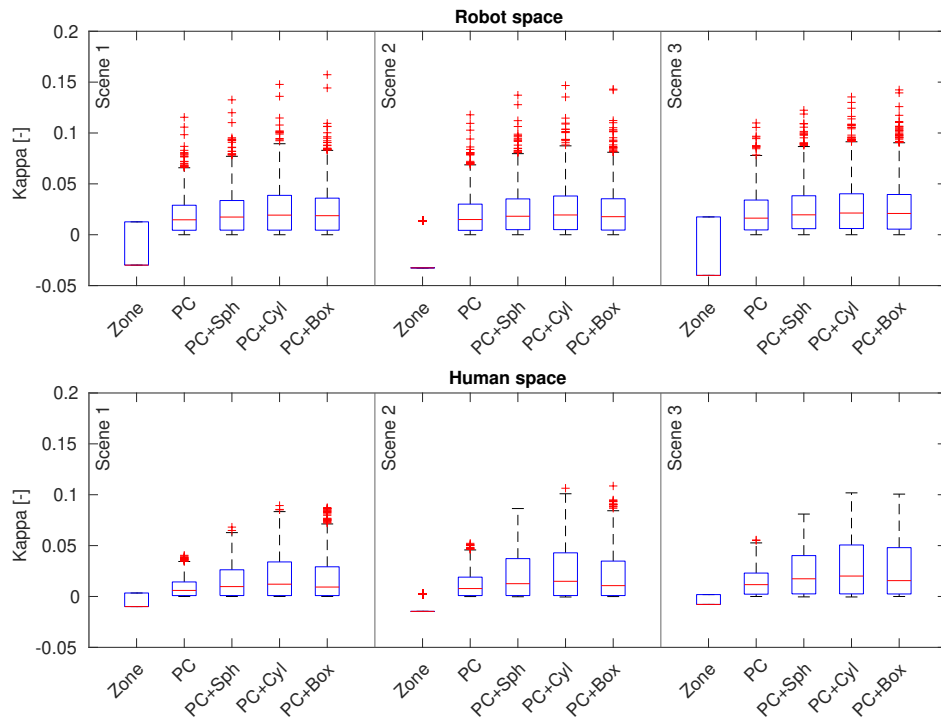


Figure 6.7: κ results in three static scenes for different interpretations of data from RGB-D camera—2D keypoint detection (Zone); raw 3D point cloud (PC); raw 3D point cloud with added spheres (PC+Sph), Raw 3D point cloud with added cylinders (PC+Cyl); raw 3D point cloud with added bounding box (PC+Box). Originally published in [129].




Figure 6.8: RGB-D camera placement heatmap for ‘robot’ space in Scene 3; F_1 (a) and κ (b) values. Originally published in [129].




Figure 6.9: Point clouds for the RGB-D camera poses with the highest metrics value for ‘robot’ space; x-axis coordinate is color-coded for better visibility. Originally published in [129].

Exp. (Scenes)	Sensor (# of poses)	Parameter	Value
1 (1,2,3)	RGB(-D) camera (172 × 5 ori)	FoV (hor × ver) [°]	87 × 58
		Res. (hor × ver) [px]	1280 × 720
		Range (RGBD only) [m]	0.6 - 6
2 (1,2,3)	Pressure pad (24)	Dimensions [m]	1.0 × 0.75
	Robot proximity cover	Inflation [m]	0.1/0.2/0.3
	RGB-D camera (9)	FoV (hor × ver) [°] Res. (hor × ver) [px] Range [m]	87 × 58 1280 × 720 0.3 - 3
3 (Dynamic)	LIDAR sensor (172)	FoV (hor × ver) [°]	360 × 45
		Range [m]	0.5 - 20
		Ang. res. (hor / ver) [°]	0.7 / 0.7

Table 6.2: Experiment parameters. Originally published in [129].

integrated the known pose of the robot into the occupied representation to show that both possibilities are available.

In this scenario, we compare only the maximum F_1 and κ as we look for the best placement only for the combination of sensors. Figure 6.10 shows the maxima of the F_1 and κ metrics for the sensor combinations relative to the triplet values.

From the F-score point of view, we can see that the combination of all three sensors does not outperform the others remarkably. Moreover, in the case of the evaluation of the ‘human’ space, the F_1 values are almost the same for all combinations containing the pressure pad. Interestingly, the pair without the pressure pad is the second-best combination for the ‘robot’ space, followed by the combination of the pressure pad and the robot proximity cover. The F_1 results show the difference between the covered spaces by individual sensors, e.g., the pressure pad is more suitable for the ‘human’ space than two other sensors and even their combination.

For the κ results, the trends are very similar to those of the F_1 results. However, the relative differences between combinations are larger, mainly for the ‘robot’ space. The triplet is followed by the pairs containing the pressure pad and the pressure pad alone. These results suggest that using a pair of sensors instead of all three sensors can have a better price-performance ratio.

Figure 6.11 shows the occupancy voxels for the triplet, the pressure pad with the robot proximity cover, and the pressure pad with the RGB-D camera combinations. For each combination, the variant with the highest ‘human’ space values of F_1 and κ for Scene 2 is shown.

2.3 Dynamic scene experiment (Exp 3)

In the third experiment, we look for the best position of the LIDAR sensor in the dynamic scene represented by 27 snapshots of the scene where the human and the robot were moving. We computed the coverage of the LIDAR sensor in the same 172 positions as those used in Exp 1.

In this experiment, we evaluate only the LIDAR placement for the ‘robot’ space coverage. Similarly to the first experiment, we created a heatmap, where each camera position is represented by its highest F_1 or κ value of the five orientations. Figure 6.8 shows the heat maps for the point cloud with the added keypoint cylinders. Again, we have differences between the

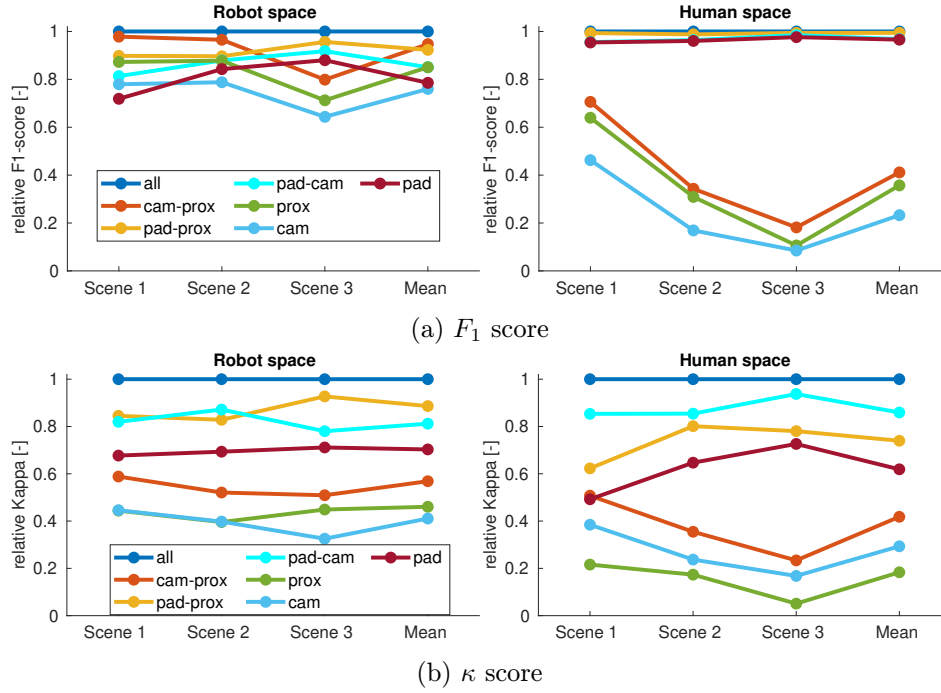



Figure 6.10: Maxima of F_1 (a) and κ (b) values in three static scenes and mean over scenes for all possible combinations of a pressure pad (pad), a robot proximity cover (prox), and an RGB-D camera (cam) sensors. Originally published in [129].

best positions based on the F_1 and κ values. The F_1 results suggest placing the LIDAR in the right half of the ceiling, which is totally different compared to the placement of the RGB-D camera in the first experiment (see Fig. 6.8). From the ceiling positions, the κ results emphasize only a few places close to the center of the ceiling, and they also propose placements on the right wall. The simulated point clouds for the LIDAR placement with the highest ‘robot’ space values of F_1 and κ are shown in Fig. 6.13 for two different snapshots. As can be seen, the best position based on the F_1 metric is on the ceiling. On the other hand, the position with the highest κ values is on the right wall (in the same notation as for the heatmap).



(a) Highest F_1 and κ for the proximity cover and the pressure pad. (b) Highest F_1 and κ for the pressure pad and the RGB-D camera. (c) Highest F_1 for the pressure pad and the RGB-D camera. (d) Highest κ for the pressure pad and the RGB-D camera.

Figure 6.11: Occupancy voxels of occupied space (blue) and free space (green) for best variants of sensor combinations for both metrics. Originally published in [129].




Figure 6.12: LIDAR sensor placement heatmap for ‘robot’ space in dynamic scene; F_1 (a) and κ (b) values. Originally published in [129].




Figure 6.13: Point clouds for the LIDAR sensor position with the highest ‘robot’ space values of F_1 (sensor on the ceiling) and κ (sensor on the wall) for the dynamic scene; x-axis coordinate is color-coded for better visibility. Originally published in [129].

■ 3 Conclusion

We introduced the definition of the perirobot space (PeRS)—the monitored region of interest for robot interaction where sensor data are represented as occupancy information. This study presented the formalization and evaluation of PeRS. We demonstrated its use in an RGB-D camera placement experiment, a multi-sensor coverage experiment, and a dynamic scene experiment. The occupancy representation allowed us to compare the effectiveness of various sensor setups and we used the well-established metrics of F_1 and κ score to evaluate the coverage of regions of interest. Therefore, our approach can serve as a prototyping tool to establish the sensor setup that provides the most efficient coverage with respect to the given metrics and sensor representations. For that reason, we made the implementation publicly available ².

■ 4 Discussion

The central idea of our approach is the simple evaluation of integrated sensors with different properties. We consider this work a first step that could be extended in multiple ways. Our approach can be interpreted in two ways. First, given a sensor setup, what are the proper, i.e., safe and efficient, ways of achieving the task. Or, given a task, determine the optimal sensor setup. For our presented multi-sensor grid search to determine the optimal sensor setup, we could consider additional criteria (e.g., space limitations or sensor costs to find the best price–performance ratio).

Our approach can capture the differences between sensors of the same type, as revealed by different suggestions for placing the LIDAR sensor and the RGB-D camera (see Figs. 6.8, 6.12). Moreover, the known robot pose can be integrated into the occupied space, as shown in the multi-sensor experiment (see Fig. 6.11), to disfavor sensor variants observing only the robot and not the space around. We found a difference between the F_1 and κ scores for the best placements for the sensors. This is probably caused by the F_1 score overlooking the correctly detected free space.

While we presented only a few sensors in this paper, our occupancy-based approach allows the easy addition of new sensors. For a detailed evaluation of a sensor setup, all appropriate representations should be considered.

The presented approach dealt with regions of interest surrounding the robot and the human. However, one could take into account the instantaneous robot speed to determine an appropriate surrounding volume for the robot to be considered as the monitored space. Or, following the peripersonal space approaches, one could ‘enact’ PeRS by training an artificial neural network on sensor data. A more industrial approach would be to integrate our occupancy-based representation with automated risk analysis as with the ADJUST tool presented in [132].

Other possible additions include additional sensors (e.g., moving sensors, sensors on humans), further representations (e.g., swept volumes for the human), taking into account the human attention (e.g., by tracking human gaze and incorporating this attention as a factor), using the occupancy representation of sensors for control (e.g., evasion of the occupied space). Also, we would like to extend the framework with a user interface to allow its easier use and optimize the code for faster computation.

²<https://github.com/ctu-vras/perirobot-space>

Conclusions and Future Work

This thesis investigated the intersection of three topics: safety, efficiency, and perception in the context of collaborative robots. The presented results show that the capabilities of collaborative robots allow more efficient close physical human-robot interaction (pHRI) than if the technical specification ISO/TS 15066 is naively interpreted. Although I refer to collaborative robots, we must always consider the whole application [22] to evaluate the safety and not just the properties of the robot. I presented my work in three thematic parts and reached the following conclusions.

In-situ measurements are indispensable for accurate risk analysis of a robot application. In Part I: Safety in Physical Human-Robot Interaction, I showed with a total of 3184 measurements of collision forces that the suggested technical specification omits certain key aspects of collisions that influence the impact forces, namely I studied with my colleagues the effect of the impact location in the robot's workspace and the effect of safety covers. I showed also how we extended the original formula with the properties of soft covers that were merely mentioned by the standard but not discussed further. In-situ measurements play a crucial role in accurately estimating the collision forces and therefore the application's safety. Directly measured impact forces are the most accurate representation of the impact forces in the robot application as they take into account also factors that cannot be modeled (e. g., robot controllers, motor behavior, or specific distribution of weight that are all proprietary). Therefore the recommendation is to always perform in-situ measurements to verify the safety of a collaborative application.

Suitable combination of safety regimes improves the efficiency of pHRI. The goal of Part II: Efficient Safe Collaboration was to show how the combination of the Power and Force Limiting regime with the Speed and Separation Monitoring regime improves the performance of the robot application while maintaining its safety. This idea to combine the regimes was shared by also other researchers in the field. Still, the three presented papers [21, 109, 110] showed each a different realization of the idea and therefore provided an exploration of the merits of the regime combination. The results of all three of the papers showed that the performance of the combined regimes is better than if the regimes are applied separately. Roughly, the combination of these regimes allows reaping the benefits of each one of them while mitigating their drawbacks. Speed and Separation Monitoring could leave the robot idle for too long, but it allows the robot to move fast if the human is not present. Power and Force Limiting regime forces the robot to move slowly but allows the robot to move all the time. I consider the merit of the presented research in providing experimental verification of this hypothesis and an easily applicable implementation of the regime combination. The presented papers used human keypoint detection to monitor the separation distance between the robot and the human keypoints. However, the shown research has presented only a possible future of human-robot interaction as the employed human keypoint detection algorithms are not yet safety-rated. Their use allows to discriminate between various parts of the human body and therefore be more sensitive to safety risks.

Sensors should be evaluated with respect to the space surrounding the robot in a formal representation. Finally, in Part III: Perception and Representation, I suggested focusing on the perirobot space, i. e., space surrounding the robot, for the evaluation of the use of effective safety measures. The presented study was merely a first step in the direction of the formalization and evaluation of perirobot space. The therein-used occupancy representation allows the comparison of various sensory setups for the robot application. The result shows that the optimal sensor setup strongly depends on the application and the applied sensors. The suggested occupancy representation manages to capture the fundamental differences between sensors. For example, it captures the differences between the various keypoint models and also the effect of combining these sensors. As opposed to classical approaches, I suggest keeping the robot at the central point of the analysis of sensor coverage and thereby include sensors that are part of the robot (e. g., sensitized skins) more naturally. The presented perirobot space concept allows the formal evaluation of sensory setups and therefore makes it easier to evaluate the suitability of sensory setups.

These three main take-away messages from my research lead to the statement that for efficient safe physical human-robot interaction a holistic yet robot-specific approach should be taken. And this holistic approach is not limited to industrial robots but should extend also to other ‘non-industrial’ robots where many of the learned lessons could be applied (as shown with the GARMi robot). This, however, asks for the continued development of certain approaches and technologies. Further investigations should be made in perirobot space models that would account for the individual properties of the robots and sensors and present possibly other than occupancy-based representations. The promise of a perirobot space approach is also its automation for the design and evaluation of sensory setups. The combination of an in-situ created collision force map together with the perirobot space sensor evaluation would help the design of safe robot applications. New factors should be also modeled and taken into account to make the approach truly holistic, for example, ‘soft’ aspects of human-robot interaction, i. e., the human interpretation of robot movements as in the Expectable Motion Unit [133]. Additionally, new safety-rated sensors should be developed and introduced that would allow the most efficient, i. e., human keypoint-based, approaches to safety regime combination leave laboratories (e. g., the camera from the Smart Robots [30] or the SICK sensors [29]). Last, but not least, the research presented in Part I: Safety joined the current attempts to improve the ISO/TS 15066 standard. With these steps, the lessons learned from safety research with collaborative robots, safe and efficient close physical human-robot collaboration could become common not only in laboratories but also in factories and our homes.

Publications

Main publications created during my studies listed in their reversed chronological order within categories. Each publication is followed by author share and CRediT authorship contribution statement. The articles are grouped according to their impact and relevance into separate sections.

Thesis subject-related publications

Impacted journal publications

1. Svarny, P.; Rozlivek, J.; Rustler, L.; Sramek, M.; Deli, Ö.; Zillich, M.; Hoffmann, M.: Effect of Active and Passive Protective Soft Skins on Collision Forces in Human–Robot Collaboration. In: *Robotics and Computer-Integrated Manufacturing*, 78, 2022, p. 102363.
Svarny, P. (20 %): Conceptualization, Methodology, Writing—original draft, Visualization, Project administration.
Rozlivek, J. (20 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.
Rustler, L. (20 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.
Sramek, M. (10 %): Methodology, Software, Investigation.
Deli, Ö. (10 %): Software, Investigation.
Zillich, M. (10 %): Conceptualization, Methodology, Writing—review & editing.
Hoffmann, M. (10 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.

Conference proceedings (rank A, B)

1. Svarny, P.; Rozlivek, J.; Rustler, L.; Hoffmann, M.: 3D Collision-Force-Map for Safe Human-Robot Collaboration. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3829–3835.
Svarny, P. (25 %): Conceptualization, Methodology, Writing—original draft, Visualization, Project administration.
Rozlivek, J. (25 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.
Rustler, L. (25 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.
Hoffmann, M. (25 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.
2. Svarny, P.; Tesar, M.; Behrens, J. K.; Hoffmann, M.: Safe Physical HRI: Toward a Unified Treatment of Speed and Separation Monitoring together with Power and Force

Limiting In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019)*. IEEE, 2019, pp. 7580–7587.

Svarny, P. (40 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization, Project administration.

Tesar, M. (20 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft.

Behrens, J. K. (20 %): Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Hoffmann, M. (20 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.

Other

1. Svarny, P.; Hamad, M.; Kurdas, A.; Hoffmann, M.; Abdolshah, S.; Haddadin, S.: Functional Mode Switching for Safe and Efficient Human-Robot Interaction In: *2022 IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2022, pp. 888–894.

Svarny, P. (25 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing—original draft, Project administration, Visualization.

Hamad, M. (25 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Kurdas, A. (20 %): Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Hoffmann, M. (10 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.

Abdolshah, S. (10 %): Writing—review & editing, Supervision, Project administration.

Haddadin, S. (10 %): Writing—review & editing, Supervision, Funding acquisition.

2. Zardykhan, D.; Svarny P.; Hoffmann M.; Shahriari E.; Haddadin S.: Collision Preventing Phase-Progress Control for Velocity Adaptation in Physical Human-Robot Interaction In: *2019 IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2019, pp. 266–273.

Zardykhan, D. (35 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Svarny, P. (20 %): Software, Investigation, Writing—original draft.

Hoffmann, M. (15 %): Writing—review & editing, Funding acquisition.

Shahriari, E. (20 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Haddadin, S. (10 %): Writing—review & editing, Supervision, Funding acquisition.

3. Svarny, P.; Straka, Z.; Hoffmann, M.: Versatile Distance Measurement between Robot and Human Key Points Using RGB-D Sensors for Safe HRI In: *1st Workshop on Proximity Perception in Robotics at IROS 2018*. KIT Scientific Publishing, 2018.

Svarny, P. (70 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization, Project administration.

Straka, Z. (10 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Hoffmann, M. (20 %): Conceptualization, Methodology, Writing—review & editing,

Supervision, Project administration, Funding acquisition.

4. Svarny, P.; Straka, Z.; Hoffmann, M.: Toward Safe Separation Distance Monitoring from RGB-D Sensors in Human-Robot Interaction In: *Proceedings of the International PhD Conference on Safe and Social Robots*. Strasbourg: Commission of the European Communities, 2018, pp. 11-14.

Svarny, P. (70 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization, Project administration.

Straka, Z. (10 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Hoffmann, M. (20 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.

Submitted publications

1. Rozlivek, J.; Svarny, P.; Hoffmann, M.: Perirobot space representation for HRI: measuring and designing collaborative workspace coverage by diverse sensors In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023)*.

Rozlivek, J. (40 %): Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing—original draft, Visualization.

Svarny, P. (40 %): Conceptualization, Methodology, Writing—original draft, Visualization, Project administration.

Hoffmann, M. (20 %): Conceptualization, Methodology, Writing—review & editing, Supervision, Project administration, Funding acquisition.

■ Publications not related to thesis subject

1. Lehmann, H.; Svarny, P.: Using a Social Robot for Different Types of Feedback during University Lectures. In: *Education Sciences & Society-Open Access 12, no. 2*, 2021.

Lehmann, H. (70 %): Conceptualization, Methodology, Writing—original draft, Formal analysis, Visualization, Project administration.

Svarny, P. (30 %): Software, Writing—review & editing.

Bibliography

- [1] S. Haddadin and E. Croft, “Physical Human–robot Interaction,” in *Springer Handbook of Robotics*. Springer, 2016, pp. 1835–1874.
- [2] P. A. Lasota, T. Fong, J. A. Shah *et al.*, *A Survey of Methods for Safe Human–robot Interaction*. Now Publishers Delft, the Netherlands, 2017, vol. 104.
- [3] F. Vicentini, “Collaborative Robotics: A Survey,” *Journal of Mechanical Design*, vol. 143, no. 4, 2021.
- [4] ISO, “ISO 8373 Robots and Robotic Devices – Vocabulary,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2012.
- [5] J. Bessler, L. Schaake, C. Bidard, J. H. Buurke, A. E. Lassen, K. Nielsen, J. Saenz, and F. Vicentini, “COVR–Towards Simplified Evaluation and Validation of Collaborative Robotics Applications Across a Wide Range of Domains Based on Robot Safety Skills,” in *International Symposium on Wearable Robotics*. Springer, 2018, pp. 123–126.
- [6] L. Gualtieri, E. Rauch, and R. Vidoni, “Emerging Research Fields in Safety and Ergonomics in Industrial Collaborative Robotics: A Systematic Literature Review,” *Robotics and Computer-Integrated Manufacturing*, vol. 67, p. 101998, 2021.
- [7] S. Haddadin, A. De Luca, and A. Albu-Schäffer, “Robot Collisions: A Survey on Detection, Isolation, and Identification,” *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.
- [8] ISO, “ISO 13849 Safety of Machinery — Safety-Related Parts of Control Systems,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2015.
- [9] —, “Safety of Machinery — Emergency Stop Function — Principles for Design,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2015.
- [10] —, “ISO 10218 Robots and Robotic Devices – Safety Requirements for Industrial Robots,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2011.
- [11] —, “ISO/TS 15066 Robots and Robotic Devices – Collaborative Robots,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2016.
- [12] P. Švarný and M. Hoffmann, “Jak číst standard(y) a něco si z toho vzít,” in *Kognícia a umelý život 2019*. Univerzita KOMenského V BRatislave, 2019, pp. 116–118.
- [13] E. Fosch Villaronga, “ISO 13482: 2014 and Its Confusing Categories. Building A Bridge Between Law and Robotics,” in *New Trends in Medical and Service Robots*. Springer, 2016, pp. 31–44.

- [14] T. Jacobs and G. S. Virk, “ISO 13482 - the New Safety Standard for Personal Care Robots,” in *ISR/Robotik 2014; 41st International Symposium on Robotics*. VDE, 2014, pp. 1–6.
- [15] A. K. Pandey and R. Gelin, “A Mass-Produced Sociable Humanoid Robot: Pepper: the First Machine of Its Kind,” *IEEE Robotics & Automation Magazine*, vol. 25, no. 3, pp. 40–48, 2018.
- [16] M. Miyagawa, Y. Kai, Y. Yasuhara, H. Ito, F. Betriana, T. Tanioka, R. Locsin *et al.*, “Consideration of Safety Management When Using Pepper, a Humanoid Robot for Care of Older Adults,” *Intelligent Control and Automation*, vol. 11, no. 01, p. 15, 2019.
- [17] ISO, “ISO 13482 Robots and Robotic Devices—Safety Requirements for Personal Care Robots,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2014.
- [18] Industrial Truck Standards Development Foundation. (2014) AGV standards. [Online]. Available: <http://www.Itsdf.Org/Cue/B56-Standards.Html>
- [19] M. Tröbinger, C. Jähne, Z. Qu, J. Elsner, A. Reindl, S. Getz, T. Goll, B. Loinger, T. Loibl, C. Kugler *et al.*, “Introducing GARMi-A Service Robotics Platform to Support the Elderly at Home: Design Philosophy, System Overview and First Results,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5857–5864, 2021.
- [20] MIRMI - Robotics and Machine Intelligence, “GARMi Concept - Robot Assistant for the Elderly,” 2019.
- [21] P. Svarny, M. Hamad, A. Kurdas, M. Hoffmann, S. Abdolshah, and S. Haddadin, “Functional Mode Switching for Safe and Efficient Human–robot Interaction,” in *2022 IEEE-RAS 21th International Conference on Humanoid Robots (Humanoids)*, 2022, pp. 888–894.
- [22] F. Vicentini, “Terminology in Safety of Collaborative Robotics,” *Robotics and Computer-Integrated Manufacturing*, vol. 63, p. 101921, 2020.
- [23] D. Henrich and S. Kuhn, “Modeling Intuitive Behavior for Safe Human/Robot Coexistence Cooperation,” in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE, 2006, pp. 3929–3934.
- [24] F. Flacco and A. De Luca, “Safe Physical Human–robot Collaboration,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 2072–2072.
- [25] M. Fujii, H. Murakami, and M. Sonehara, “Study on Application of a Human–robot Collaborative System Using Hand-Guiding in a Production Line,” *IHI Eng. Rev*, vol. 49, no. 1, pp. 24–29, 2016.
- [26] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh, “Realtime Multi-Person 2d Pose Estimation Using Part Affinity Fields,” in *CVPR*, vol. 1, no. 2, 2017, p. 7.
- [27] E. Insafutdinov, L. Pishchulin, B. Andres, M. Andriluka, and B. Schiele, “Deepcut: A Deeper, Stronger, and Faster Multi-Person Pose Estimation Model,” in *European Conference on Computer Vision*. Springer, 2016, pp. 34–50.

- [28] Pilz Automation Safety L.P., “Pilz SafetyEye Announcement,” 2008. [Online]. Available: <https://news.thomasnet.Com/fullstory/safe-camera-system-offers-3d-safety-monitoring-and-control-540923>
- [29] Sick, “Sick Product Page,” 2022. [Online]. Available: <https://www.sick.com/pl/en/safety-camera-sensors/safety-camera-sensors/safevisionary2/>
- [30] Smart Robots, “Smart Robots,” 2022. [Online]. Available: <http://www.smartrobots.it/>
- [31] B. Lacevic and P. Rocco, “Kinetostatic Danger Field-A Novel Safety Assessment for Human–robot Interaction,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2010, pp. 2169–2174.
- [32] O. Khatib, “Real-Time Obstacle Avoidance for Manipulators and Mobile Robots,” *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [33] F. Flacco, T. Kroeger, A. De Luca, and O. Khatib, “A Depth Space Approach for Evaluating Distance to Objects,” *Journal of Intelligent & Robotic Systems (JINS)*, vol. 80, p. 7, 2015.
- [34] F. Fabrizio and A. De Luca, “Real-Time Computation of Distance to Dynamic Obstacles with Multiple Depth Sensors,” *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 56–63, 2016.
- [35] J. A. Marvel, “Performance Metrics of Speed and Separation Monitoring in Shared Workspaces,” *IEEE Transactions on Automation Science and Engineering*, vol. 10, no. 2, pp. 405–414, 2013.
- [36] J. A. Marvel and R. Norcross, “Implementing Speed and Separation Monitoring in Collaborative Robot Workcells,” *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 144–155, 2017.
- [37] A. M. Zanchettin, N. M. Ceriani, P. Rocco, H. Ding, and B. Matthias, “Safety in Human–robot Collaborative Manufacturing Environments: Metrics and Control,” *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 2, pp. 882–893, 2016.
- [38] A. M. Zanchettin, A. Casalino, L. Piroddi, and P. Rocco, “Prediction of Human Activity Patterns for Human–robot Collaborative Assembly Tasks,” *IEEE Transactions on Industrial Informatics*, vol. 15, no. 7, pp. 3934–3942, 2018.
- [39] A. M. Zanchettin, P. Rocco, S. Chiappa, and R. Rossi, “Towards An Optimal Avoidance Strategy for Collaborative Robots,” *Robotics and Computer-Integrated Manufacturing*, vol. 59, pp. 47–55, 2019.
- [40] B. Lacevic, A. M. Zanchettin, and P. Rocco, “Safe Human-Robot Collaboration via Collision Checking and Explicit Representation of Danger Zones,” *IEEE Transactions on Automation Science and Engineering*, 2022.
- [41] S. Haddadin, M. Suppa, S. Fuchs, T. Bodenmüller, A. Albu-Schäffer, and G. Hirzinger, “Towards the Robotic Co-Worker,” in *Robotics Research*. Springer, 2011, pp. 261–282.
- [42] P. Fiorini and Z. Shiller, “Motion Planning in Dynamic Environments Using Velocity Obstacles,” *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, 1998.

- [43] R. Vatcha and J. Xiao, “Perceiving Guaranteed Continuously Collision-Free Robot Trajectories in An Unknown and Unpredictable Environment,” in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference On*. IEEE, 2009, pp. 1433–1438.
- [44] A. M. Zanchettin and B. Lavecic, “Safe and Minimum-time Path-following Problem for Collaborative Industrial Robots,” *Journal of Manufacturing Systems*, vol. 65, pp. 686–693, 2022.
- [45] A. De Luca and F. Flacco, “Integrated Control for PHRI: Collision Avoidance, Detection, Reaction and Collaboration,” in *Biomedical Robotics and Biomechatronics (BioRob), 2012 4th IEEE RAS & EMBS International Conference On*. IEEE, 2012, pp. 288–295.
- [46] F. Flacco, T. Kröger, A. De Luca, and O. Khatib, “A Depth Space Approach to Human–robot Collision Avoidance,” in *2012 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2012, pp. 338–345.
- [47] D. H. P. Nguyen, M. Hoffmann, A. Roncone, U. Pattacini, and G. Metta, “Compact Real-Time Avoidance on a Humanoid Robot for Human–robot Interaction,” in *Proceedings of the 2018 ACM/IEEE International Conference on Human–robot Interaction*. ACM, 2018, pp. 416–424.
- [48] C. Liu and M. Tomizuka, “Algorithmic Safety Measures for Intelligent Industrial Co-Robots,” in *Robotics and Automation (ICRA), 2016 IEEE International Conference On*. IEEE, 2016, pp. 3095–3102.
- [49] R. Weitschat and H. Aschemann, “Safe and Efficient Human–robot Collaboration Part II: Optimal Generalized Human-In-The-Loop Real-Time Motion Generation,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 3781–3788, 2018.
- [50] A. Palleschi, M. Hamad, S. Abdolshah, M. Garabini, S. Haddadin, and L. Pallottino, “Fast and Safe Trajectory Planning: Solving the Cobot Performance/Safety Trade-Off in Human–robot Shared Environments,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5445–5452, 2021.
- [51] K. Suita, Y. Yamada, N. Tsuchida, K. Imai, H. Ikeda, and N. Sugimoto, “A Failure-To-Safety” Kyozon” System with Simple Contact Detection and Stop Capabilities for Safe Human-Autonomous Robot Coexistence,” in *Robotics and Automation, International Conference On*, vol. 3. IEEE, 1995, pp. 3089–3096.
- [52] N. Mansfeld, M. Hamad, M. Becker, A. G. Marin, and S. Haddadin, “Safety Map: A Unified Representation for Biomechanics Impact Data and Robot Instantaneous Dynamic Properties,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1880–1887, 2018.
- [53] D. Han, M. Y. Park, H. Shin, K. S. Kim, and S. Rhim, “Identifying Safety Conditions of Human–robot Collision Based on Skin Injury Analysis,” in *International Conference on Ubiquitous Robots (UR)*. IEEE, 2018, pp. 420–423.
- [54] M. Y. Park, D. Han, J. H. Lim, M. K. Shin, Y. R. Han, D. H. Kim, S. Rhim, and K. S. Kim, “Assessment of Pressure Pain Thresholds in Collisions with Collaborative Robots,” *PLoS ONE*, vol. 14, no. 5, 2019.

- [55] N. Lucci, B. Lacevic, A. M. Zanchettin, and P. Rocco, “Combining Speed and Separation Monitoring with Power and Force Limiting for Safe Collaborative Robotics Applications,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6121–6128, 2020.
- [56] Z. Bi, C. Luo, Z. Miao, B. Zhang, W. Zhang, and L. Wang, “Safety Assurance Mechanisms of Collaborative Robotic Systems in Manufacturing,” *Robotics and Computer-Integrated Manufacturing*, vol. 67, p. 102022, 2021.
- [57] V. Gopinath, K. Johansen, M. Derelov, A. Gustafsson, and S. Axelsson, “Safe Collaborative Assembly on a Continuously Moving Line with Large Industrial Robots,” *Robotics and Computer-Integrated Manufacturing*, vol. 67, p. 102048, 2021.
- [58] S. Haddadin, “Physical Safety in Robotics,” in *Formal Modeling and Verification of Cyber-Physical Systems*. Springer, 2015, pp. 249–271.
- [59] P. Svarny, J. Rozlivek, L. Rustler, M. Sramek, Ö. Deli, M. Zillich, and M. Hoffmann, “Effect of Active and Passive Protective Soft Skins on Collision Forces in Human–robot Collaboration,” *Robotics and Computer-Integrated Manufacturing*, vol. 78, p. 102363, 2022.
- [60] M. Hamad, A. Kurdas, S. Abdolshah, and S. Haddadin, “Experimental Injury Biomechanics of Human Body Upper Extremities: Anatomy, Injury Severity Classification, and Impact Testing Setups,” in *2021 IEEE International Conference on Intelligence and Safety for Robotics (ISR)*. IEEE, 2021, pp. 310–315.
- [61] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, “Safety Evaluation of Physical Human–robot Interaction Via Crash-Testing,” in *Robotics: Science and Systems*, vol. 3, 2007, pp. 217–224.
- [62] —, “Dummy Crash-Tests for the Evaluation of Rigid Human–robot Impacts,” in *International Workshop on Technical Challenges for Dependable Robots in Human Environments*. Citeseer, 2007.
- [63] O. Khatib, “Inertial Properties in Robotic Manipulation: An Object-Level Framework,” *The International Journal of Robotics Research*, vol. 14, no. 1, pp. 19–36, 1995.
- [64] S.-D. Lee, B.-S. Kim, and J.-B. Song, “Human–robot Collision Model with Effective Mass and Manipulability for Design of a Spatial Manipulator,” *Advanced Robotics*, vol. 27, no. 3, pp. 189–198, 2013.
- [65] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*. Springer Science & Business Media, 2010.
- [66] S. Haddadin, A. Albu-Schäffer, A. De Luca, and G. Hirzinger, “Collision Detection and Reaction: A Contribution to Safe Physical Human–robot Interaction,” in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008, pp. 3356–3363.
- [67] S. Haddadin, S. Haddadin, A. Khoury, and Rokahr, “On Making Robots Understand Safety: Embedding Injury Knowledge into Control,” *The International Journal of Robotics Research (IJRR)*, vol. 31, no. 13, pp. 1578–1602, 2012.

- [68] H. Shin, K. Seo, and S. Rhim, “Allowable Maximum Safe Velocity Control Based on Human–robot Distance for Collaborative Robot,” in *Int. Conf. on Ubiquitous Robots (UR)*. IEEE, 2018, pp. 401–405.
- [69] S. Haddadin, A. Albu-Schaffer, and G. Hirzinger, “The Role of the Robot Mass and Velocity in Physical Human–robot Interaction-Part I: Non-Constrained Blunt Impacts,” in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1331–1338.
- [70] —, “The Role of the Robot Mass and Velocity in Physical Human–robot Interaction-Part II: Constrained Blunt Impacts,” in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1339–1345.
- [71] K. L. Johnson and K. L. Johnson, *Contact Mechanics*. Cambridge University Press, 1987.
- [72] B. R. Vemula, M. Ramteen, G. Spampinato, and B. Fagerström, “Human–robot Impact Model: for Safety Assessment of Collaborative Robot Design,” in *2017 IEEE International Symposium on Robotics and Intelligent Sensors (IRIS)*. IEEE, 2017, pp. 236–242.
- [73] B. Vemula, B. Matthias, and A. Ahmad, “A Design Metric for Safety Assessment of Industrial Robot Design Suitable for Power-And Force-Limited Collaborative Operation,” *International Journal of Intelligent Robotics and Applications*, vol. 2, no. 2, pp. 226–234, 2018.
- [74] N. Kovincic, H. Gattringer, A. Müller, M. Weyrer, A. Schlotzhauer, L. Kaiser, and M. Brandstötter, “A Model-Based Strategy for Safety Assessment of a Robot Arm Interacting with Humans,” *PAMM*, vol. 19, no. 1, p. E201900247, 2019.
- [75] N. Kovincic, H. Gattringer, M. Andreas, and M. Brandstötter, “A Boosted Decision Tree Approach for a Safe Human–robot Collaboration in Quasi-Static Impact Situations,” *Advances in Service and Industrial Robotics: Results of RAAD*, p. 235, 2020.
- [76] A. Schlotzhauer, L. Kaiser, J. Wachter, M. Brandstötter, and M. Hofbauer, “On the Trustability of the Safety Measures of Collaborative Robots: 2D Collision-Force-Map of a Sensitive Manipulator for Safe HRC,” in *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2019, pp. 1676–1683.
- [77] P. Aivaliotis, S. Aivaliotis, C. Gkournelos, K. Kokkalis, G. Michalos, and S. Makris, “Power and Force Limiting on Industrial Robots for Human–robot Collaboration,” *Robotics and Computer-Integrated Manufacturing*, vol. 59, pp. 346–360, 2019.
- [78] C. Sloth and H. G. Petersen, “Computation of Safe Path Velocity for Collaborative Robots,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 6142–6148.
- [79] A. Meguenani, V. Padois, and P. Bidaud, “Control of Robots Sharing Their Workspace with Humans: An Energetic Approach to Safety,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 4678–4684.
- [80] R. Rossi, M. P. Polverini, A. M. Zanchettin, and P. Rocco, “A Pre-Collision Control Strategy for Human–robot Interaction Based on Dissipated Energy in Potential Inelastic Impacts,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 26–31.

- [81] G. Cheng, E. Dean-Leon, F. Bergner, J. R. G. Olvera, Q. Leboutet, and P. Mitterdorfer, “A Comprehensive Realization of Robot Skin: Sensors, Sensing, Control, and Applications,” *Proceedings of the IEEE*, vol. 107, no. 10, pp. 2034–2051, 2019.
- [82] S. Tsuji and T. Kohama, “Self-Capacitance Proximity and Tactile Skin Sensor with Shock-Absorbing Structure for a Collaborative Robot,” *IEEE Sensors Journal*, 2020.
- [83] M. Bdiwi, M. Pfeifer, and A. Sterzing, “A New Strategy for Ensuring Human Safety During Various Levels of Interaction with Industrial Robots,” *CIRP Annals*, vol. 66, no. 1, pp. 453–456, 2017.
- [84] F. Ferraguti, M. Bertuletti, C. T. Landi, M. Bonfe, C. Fantuzzi, and C. Secchi, “A Control Barrier Function Approach for Maximizing Performance While Fulfilling to ISO/TS 15066 Regulations,” *IEEE Robotics and Automation Letters*, 2020.
- [85] R. J. Kirschner, A. Kurdas, K. Karacan, P. Junge, S. A. B. Birjandi, N. Mansfeld, S. Abdolshah, and S. Haddadin, “Towards A Reference Framework for Tactile Robot Performance and Safety Benchmarking,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4290–4297.
- [86] V. Di Cosmo, A. Giusti, R. Vidoni, M. Riedl, and D. T. Matt, “Collaborative Robotics Safety Control Application Using Dynamic Safety Zones Based on the ISO/TS 15066: 2016,” in *International Conference on Robotics in Alpe-Adria Danube Region*. Springer, 2019, pp. 430–437.
- [87] C. Byner, B. Matthias, and H. Ding, “Dynamic Speed and Separation Monitoring for Collaborative Robot Applications—Concepts and Performance,” *Robotics and Computer-Integrated Manufacturing*, vol. 58, pp. 239–252, 2019.
- [88] M. Safeea and P. Neto, “Minimum Distance Calculation Using Laser Scanner and IMUs for Safe Human–robot Interaction,” *Robotics and Computer-Integrated Manufacturing*, vol. 58, pp. 33–42, 2019.
- [89] B. Wang, “Coverage Problems in Sensor Networks: A Survey,” *ACM Computing Surveys (CSUR)*, vol. 43, no. 4, pp. 1–53, 2011.
- [90] R. Elhabyan, W. Shi, and M. St-Hilaire, “Coverage Protocols for Wireless Sensor Networks: Review and Future Directions,” *Journal of Communications and Networks*, vol. 21, no. 1, pp. 45–60, 2019.
- [91] P. Oščádal, T. Kot, T. Spurný, J. Suder, M. Vocetka, L. Dobeš, and Z. Bobovský, “Camera Arrangement Optimization for Workspace Monitoring in Human–robot Collaboration,” *Sensors*, vol. 23, no. 1, p. 295, 2023.
- [92] M. B. Alatise and G. P. Hancke, “A Review on Challenges of Autonomous Mobile Robot and Sensor Fusion Methods,” *IEEE Access*, vol. 8, pp. 39 830–39 846, 2020.
- [93] V. Magnanimo, S. Walther, L. Tecchia, C. Natale, and T. Guhl, “Safeguarding a Mobile Manipulator Using Dynamic Safety Fields,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 2972–2977.
- [94] A. Rashid, K. Peesapati, M. Bdiwi, S. Krusche, W. Hardt, and M. Putz, “Local and Global Sensors for Collision Avoidance,” in *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2020, pp. 354–359.

- [95] D. Susanj, D. Pincic, and K. Lenac, “Effective Area Coverage of 2D and 3D Environments with Directional and Isotropic Sensors,” *IEEE Access*, vol. 8, pp. 185 595–185 608, 2020.
- [96] K. He, R. Newbury, T. Tran, J. Haviland, B. Burgess-Limerick, D. Kulić, P. Corke, and A. Cosgun, “Visibility Maximization Controller for Robotic Manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8479–8486, 2022.
- [97] I. Ibrahim, F. Farshidian, J. Preisig, P. Franklin, P. Rocco, and M. Hutter, “Whole-Body MPC and Dynamic Occlusion Avoidance: A Maximum Likelihood Visibility Approach,” pp. 221–227, 2022.
- [98] M. N. Finean, W. Merkt, and I. Havoutis, “Where Should I Look? Optimized Gaze Control for Whole-Body Collision Avoidance in Dynamic Environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1095–1102, 2021.
- [99] M. P. Polverini, A. M. Zanchettin, and P. Rocco, “A Computationally Efficient Safety Assessment for Collaborative Robotics Applications,” *Robotics and Computer-Integrated Manufacturing*, vol. 46, pp. 25–37, 2017.
- [100] A. Roncone, M. Hoffmann, U. Pattacini, L. Fadiga, and G. Metta, “Peripersonal Space and Margin of Safety Around the Body: Learning Tactile-Visual Associations in a Humanoid Robot with Artificial Skin,” *PLoS ONE*, vol. 11, no. 10, p. E0163713, 2016.
- [101] L. Scalera, R. Vidoni, and A. Giusti, “Optimal Scaling of Dynamic Safety Zones for Collaborative Robotics,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 3822–3828.
- [102] R.-J. Halme, M. Lanz, J. Kämäräinen, R. Pieters, J. Latokartano, and A. Hietanen, “Review of Vision-Based Safety Systems for Human–robot Collaboration,” *Procedia CIRP*, vol. 72, pp. 111–116, 2018.
- [103] S. Kumar, S. Arora, and F. Sahin, “Speed and Separation Monitoring Using On-Robot Time-Of-Flight Laser-Ranging Sensor Arrays,” in *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2019, pp. 1684–1691.
- [104] E. Dean, K. R. Amaro, F. Bergner, I. Dianov, and G. Cheng, “Integration of Robotic Technologies for Rapidly Deployable Robots,” *IEEE Transactions on Industrial Informatics*, 2017.
- [105] K. Qi, Z. Song, and J. S. Dai, “Safe Physical Human–robot Interaction: A Quasi Whole-Body Sensing Method Based on Novel Laser-Ranging Sensor Ring Pairs,” *Robotics and Computer-Integrated Manufacturing*, vol. 75, p. 102280, 2022.
- [106] S. Savazzi, V. Rampa, F. Vicentini, and M. Giussani, “Device-Free Human Sensing and Localization in Collaborative Human–robot Workspaces: A Case Study,” *IEEE Sensors Journal*, vol. 16, no. 5, pp. 1253–1264, 2016.
- [107] S. E. Navarro, S. Mühlbacher-Karrer, H. Alagi, H. Zangl, K. Koyama, B. Hein, C. Duriez, and J. R. Smith, “Proximity Perception in Human-centered Robotics: A Survey on Sensing Systems and Applications,” *IEEE Transactions on Robotics*, vol. 38, no. 3, pp. 1599–1620, 2021.

- [108] P. Svarny, J. Rozlivek, L. Rustler, and M. Hoffmann, “3D Collision-Force-Map for Safe Human–robot Collaboration,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3829–3835.
- [109] P. Svarny, Z. Straka, and M. Hoffmann, “Toward Safe Separation Distance Monitoring From RGB-D Sensors in Human–robot Interaction,” in *International PhD Conference on Safe and Social Robotics (SSR-2018)*, 2018, pp. 11–14.
- [110] P. Svarny, M. Tesar, J. K. Behrens, and M. Hoffmann, “Safe Physical HRI: Toward a Unified Treatment of Speed and Separation Monitoring Together with Power and Force Limiting,” in *2019 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2019, pp. 7580–7587.
- [111] D. Zardykhan, P. Svarny, M. Hoffmann, E. Shahriari, and S. Haddadin, “Collision Preventing Phase-Progress Control for Velocity Adaptation in Human–robot Collaboration,” in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2019, pp. 266–273.
- [112] D. Mewes and F. Mauser, “Safeguarding Crushing Points by Limitation of Forces,” *International Journal of Occupational Safety and Ergonomics*, vol. 9, no. 2, pp. 177–191, 2003.
- [113] P. Svarny, J. Rozlivek, L. Rustler, and M. Hoffmann, “3D Collision Force Map – Public Dataset,” 2020. [Online]. Available: <https://Bit.Ly/34D27B8>
- [114] ISO, “ISO 12100 Safety of Machinery – General Principles for Design – Risk Assessment and Risk Reduction,” International Organization for Standardization, Geneva, CH, Tech. Rep., 2010.
- [115] IEC, “IEC 60204-1,” International Electrotechnical Commission, Tech. Rep., 2016.
- [116] Universal Robots, “Universal Robots - Safety FAQ,” 2020. [Online]. Available: <https://www.universal-robots.com/articles/ur/safety-faq/>
- [117] K. Hamm, Ed., *Biomechanics of Human Movement*. OpenStax, 2020. [Online]. Available: <https://Pressbooks.Bccampus.Ca/Humanbiomechanics/Chapter/Safety-Technology/>
- [118] M. Zillich and W. Feiten, “A Versatile Tactile Sensor System for Covering Large and Curved Surface Areas,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 20–24.
- [119] Intel, “Librealsense,” 2018.
- [120] G. Bradski, “The OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.
- [121] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh, “OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields,” 2019.
- [122] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh, “OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [123] D. Coleman, I. Sukan, S. Chitta, and N. Correll, “Reducing the Barrier to Entry of Complex Robotic Software: A MoveIt! Case Study,” *JOSER - Journal of Software Engineering for Robotics*, vol. 5, pp. 3–16, 2014.

- [124] V. G. Rao and D. S. Bernstein, “Naive Control of the Double Integrator,” *IEEE Control Systems Magazine*, vol. 21, no. 5, pp. 86–97, 2001.
- [125] M. Huber, M. Rickert, A. Knoll, T. Brandt, and S. Glasauer, “Human–robot Interaction in Handing-Over Tasks,” in *RO-MAN 2008-The 17th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2008, pp. 107–112.
- [126] L. Biagiotti and C. Melchiorri, *Trajectory Planning for Automatic Machines and Robots*. Springer-Verlag Berlin Heidelberg, 2008.
- [127] P. Jiménez, F. Thomas, and C. Torras, “3D Collision Detection: A Survey,” *Computers & Graphics*, vol. 25, no. 2, pp. 269–285, 2001.
- [128] R. A. Güler, N. Neverova, and I. Kokkinos, “Densepose: Dense Human Pose Estimation in the Wild,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7297–7306.
- [129] J. Rozlivek, P. Svarny, and M. Hoffmann, “Perirobot space representation for HRI: measuring and designing collaborative workspace coverage by diverse sensors,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. (submitted, arXiv:2303.02367).
- [130] B. Siciliano, O. Khatib, and T. Kröger, *Springer Handbook of Robotics*. Springer, 2008, vol. 200.
- [131] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, “OctoMap: An Efficient Probabilistic 3D Mapping Framework Based on Octrees,” *Autonomous Robots*, vol. 34, no. 3, pp. 189–206, 2013.
- [132] M. Brandstötter, T. Komenda, G. Breitenhuber, M. Rathmair, M. Steiner, C. Laflamme, A. Müller, and M. Hofbaur, “A Method to Enhance the Flexibility of Collaborative Human-robot Workspaces through an Extended Safety Perspective,” *Procedia CIRP*, vol. 112, pp. 197–202, 2022.
- [133] R. J. Kirschner, H. Mayer, L. Burr, N. Mansfeld, S. Abdolshah, and S. Haddadin, “Expectable Motion Unit: Avoiding Hazards From Human Involuntary Motions in Human–robot Interaction,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2993–3000, 2022.