Review

An atlas of physical human–robot interaction

Agostino De Santis [a], Bruno Siciliano [a,*], Alessandro De Luca [b], Antonio Bicchi [c]

[a] *Dipartimento di Informatica e Sistemistica, Università degli Studi di Napoli Federico II, Via Claudio 21, 80125 Napoli, Italy*
[b] *Dipartimento di Informatica e Sistemistica "A. Ruberti", Università degli Studi di Roma "La Sapienza", Via Eudossiana 18, 00184 Roma, Italy*
[c] *Centro Interdipartimentale di Ricerca "E. Piaggio", Università degli Studi di Pisa, Via Diotisalvi 2, 56125 Pisa, Italy*

Abstract

A broad spectrum of issues have to be addressed in order to tackle the problem of a safe and dependable physical Human–Robot Interaction (pHRI). In the immediate future, metrics related to safety and dependability have to be found in order to successfully introduce robots in everyday enviornments. While there are certainly also "cognitive" issues involved, due to the human perception of the robot (and vice versa), and other objective metrics related to fault detection and isolation, our discussion focuses on the peculiar aspects of "physical" interaction with robots. In particular, safety and dependability are the underlying evaluation criteria for mechanical design, actuation, and control architectures. Mechanical and control issues are discussed with emphasis on techniques that provide safety in an intrinsic way or by means of control components. Attention is devoted to dependability, mainly related to sensors, control architectures, and fault handling and tolerance. Suggestions are provided to draft metrics for evaluating safety and dependability in pHRI, and references to the works of the scientific groups involved in the pHRI research complete the study. The present atlas is a result of the EURON perspective research project "Physical Human–Robot Interaction in anthropic DOMains (PHRIDOM)", aimed at charting the new territory of pHRI, and constitutes the scientific basis for the ongoing STReP project "Physical Human–Robot Interaction: depENDability and Safety (PHRIENDS)", aimed at developing key components for the next generation of robots, designed to share their environment with people.
© 2007 Elsevier Ltd. All rights reserved.

Keyword: Physical human–robot interaction in anthropic domains

Contents

* Corresponding author. Tel.: +39 0817683179; fax: +39 06233226128.
  E-mail addresses: agodesa@unina.it (A. De Santis), siciliano@unina.it (B. Siciliano), deluca@dis.uniroma1.it (A. De Luca), bicchi@ing.unipi.it (A. Bicchi).

1. Introduction

The extension of application domains for robotics, from factories to human environments, is growing, due to the elderly-dominated scenario of most industrialized countries, the desire of automatizing common daily tasks, and the lack or high cost of local human expertise. Safety and dependability are the keys to a successful introduction of robots into human environments. Robots for physical assistance to humans should reduce fatigue and stress, increase human capabilities in terms of force, speed, and precision, and improve in general the quality of life; on the other hand, the human can bring experience, global knowledge, and understanding for a correct execution of tasks [1]. Only dependable robot architectures can be accepted for supporting "human-in-the-loop" conditions and human–robot teams. Application domains asking for human augmentation and substitution by robots include everyday houses and offices, but also unmanned warfare operations, mainly in USA [2], and robot companions as well as humanoids, the robots with "kokoro" (heart) diffused in Japan [3]. Moreover, teleassistance and the use of computers and devices for remote medical care pave the way to the future use of robots in domestic environments. Researchers worldwide are studying the social factors related to the introduction of robots in human environments and often their attention is focused on the cognitive interaction with machines.

Since it is impossible to model every action in an unstructured anthropic environment, the "intelligent connection of perception with action" of robots implies the presence of autonomous behavior, which is interesting per se and needed to solve real problems. However, this can result in dangerous situations for humans co-existing in the robot operational domain. When considering the current mechanical structure of robots available on the market, it is clear how physical issues are crucial, since "natural" or unexpected behavior of people during interaction with robots can result in very severe injuries.

While robots should make "independent decisions", their designers must consider physical, social and ethical implications of such autonomy [4]. In Europe, attention was recently devoted also to ethical issues [5], mainly considering issue arising for robotic/neural implants in human bodies, as discussed, e.g., in the EU project Neurobotics [6].

In order to spread the presence of robots in everyday life, personal robots just like personal computers, safety and dependability issues must be solved first. Today, computers are no more perceived as strange machines, while current robots are still heavy and unsafe. However, it must be pointed out that safety standards for pHRI are still not well defined in the scientific community. Also, efficient communication systems are crucial to have "wearable robots" analogous to "wearable" PCs. One crucial capability of a robot for pHRI is the generation of supplementary forces to overcome human physical limits [7,8]. In anthropic domains, a robot may substitute the complex infrastructure needed for "intelligent environments" [9] or telesurveillance. In these cases, instead of equipping the environment with many sensors and devices, a single robot could behave both as a sensor and an actuator, able to navigate through different rooms, sense the environment, and perform the requested task.
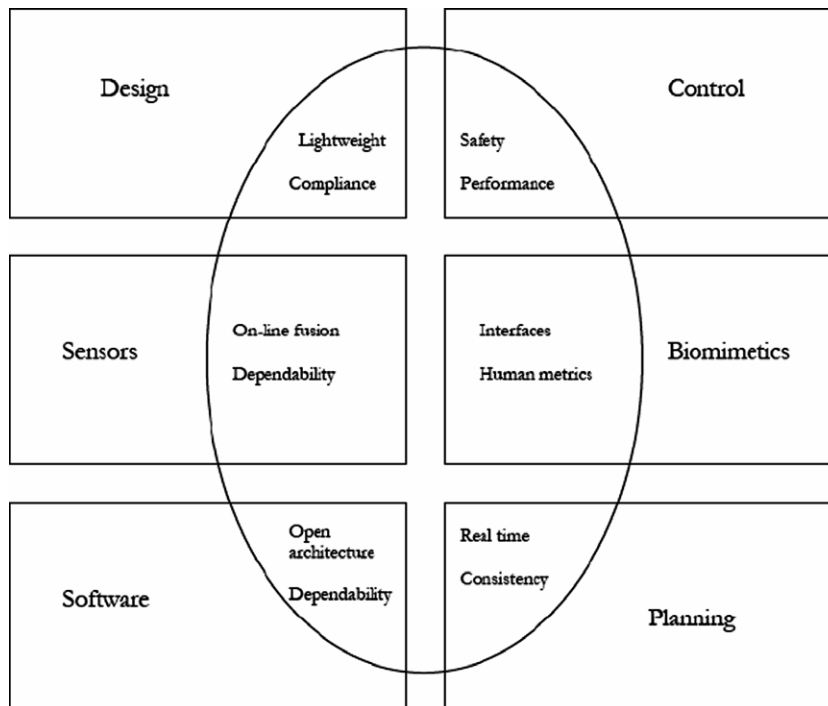
Fig. 1. A "map" of robotics for anthropic domains: main issues and superposition for pHRI.

Right now, a sort of descartes "duality" leads to accepting a dichotomy: the "brain" of robotic systems is usually studied by computer scientists and neuroscientists, whereas the study of mechanisms and their control is assigned to cybernetics, electronic, and mechanical engineers. Cognitive and physical interaction, however, are not independent: physical interaction can help in setting rules for cognitive evaluations of the environment during interaction tasks, while cognitive aspects may improve the physical interaction by setting suitable control interaction parameters. As a simple example, haptics is used to "understand" the characteristics of an environment (soft or rigid), while cognitive-based inference rules can be considered for compliance control of manipulators physically interacting with humans (if the person is a child, then the compliance should be high). Therefore, an improved analysis of the problems related to the physical interaction with robots becomes necessary. This topic must be addressed considering together the design of mechanism, sensors, actuators and control architecture in the special perspective for the interaction with humans.

When determining the differences between computers and robots from a user point of view, one should deeply consider the key problem of embodiment. People seem to perceive autonomous robots differently than they do with most other computer technologies: mental models are more anthropomorphic, and people attribute to robots human-like qualities and capabilities [10]. During a physical interaction, if human-like robotic arms are used, motion capabilities can be simpler to understand. In general, the user response to such an interactive system is always dominated by a specific mental model about how the robot behaves. If the robot looks like a living creature, the mental model of its behavior may approach the mental model of humans or of pets, and there may also be unexpected social interactions [11]. A user mental model may result in a fake robot "dependability": its zoomorphism or the human posture of a humanoid robot could give a wrong idea about the awareness of the robot, since it looks like a living creature. Indeed, mental models can be changed with experience, but anthropomorphism is still a forced consequence of our nature, especially for non-skilled users, because "our imagination cannot be anything but anthropomorphic" [12]. These mental models depend strongly on cultural approaches: a robot can be a companion or a servant. Safety issues are usually considered more relevant for robot servants, while robot companions have typically a simpler mechanical design because the focus is on cognitive interaction and not on task execution. In addition, it has been noticed in a NSF/

DARPA research project [13] that the cognitive HRI, which focuses on the combination of the user and the robot, is different with respect to a simple human–computer interaction focussing only on the person using the computer. Moreover, different roles of interaction with robots are possible since different people interact in different ways with the same robot, and the robot in turn reacts differently based on its perception of the world. Also, there are failure modes on the "robot side" that can degrade the quality of the interaction. Effective communication between a person and a robot may depend on whether there exists a common domain of understanding between the two. The interface design is crucial to let the human be aware of the robot possibilities and to provide her/him with a natural way to keep the robot under control at every time.

This "atlas" of physical human–robot interaction will present issues and discuss results from different research groups about possible metrics for the evaluation of safety, dependability, and performance in pHRI. Existing sources in this domain include workshops, European projects, and results in journal, e.g. [14–17], and conference papers. While many events in recent years were related to single aspects, our present effort is to consider together safety and dependability as the unified optimality criteria for future technical challenges in the design of robots for human environments [18]. Like in geographic atlases [19], the different sections of this document will contain some superposition in discussing such themes, which shows contact points and overlapping interests (see Fig. 1). Every paragraph is a "map" where "destinations" are machines embedded in anthropic domains, "viability conditions" are safety, dependability, reliability, failure recovery, and performance, and "via points" are sensors, actuators, mechanics, control, and software architectures. There is a need for "pathways" connecting crucial components and leading to technological solutions to applications, while fulfilling the viability requirements.

2. Safety in pHRI

2.1. General aspects on safety and dependability

In the complexity of a HRI, the physical viewpoint is mainly focused on the risks of collisions occurring between the robot and its user: too high energy/power may be transferred by the robot, resulting in serious human damages. Severity indices of injuries may be used to evaluate the safety of robots in pHRI. These should take into account the possible damages occurring when a manipulator collides with a human head, neck, chest or arm. Several standard indices of injury severity exist in other, non-robotic, domains. The automotive industry developed empirical/experimental formulas that correlate human body's acceleration to injury severity, while the suitability of such formulas is still an open issue in robotics.

To increase robot safety, all aspects of manipulator design, including mechanics, electronics, and software, should be considered. In the mechanical design, the elimination of sharp edges can reduce the potential for lacerations. The main solution for reducing the instantaneous severity of impacts is to pursue a mechanical design that reduces manipulator link inertia and weight by using lightweight but stiff materials, complemented by the presence of compliant components in the structure. Compliance can be introduced at the contact point by a soft covering of the whole arm with visco-elastic materials or by adopting compliant transmissions at the robot joints. The latter allow the actuators' rotor inertia to be dynamically decoupled from the links, whenever an impact occurs. Increasing in this way the robot mechanical compliance while reducing its overall apparent inertia has been realized though different elastic actuation/transmission arrangements which include: relocation of actuators close to the robot base and transmission of motion through steel cables and pulleys, combination of harmonic drives and lightweight link design [20], and use of parallel and distributed macro-mini actuation with elastic couplings. Other improvements for anticipating and reacting to collisions can be achieved through the use of combinations of external/internal robot sensing, electronic hardware and software safety procedures, which intelligently monitor, supervise, and control manipulator operation.

Indeed, the problem of blending the requirements for safety while keeping "traditional" robot performance (speed and accuracy) high remains an open challenge for the designers of human-centered robotic manipulators [21]. As clear from the previous discussion, safety metrics that consider only the "brain" of the robot as an intelligent machine are inappropriate in this context. Asimov's famous "three laws of robotics" are mainly science-fiction, since the "will" of the robots cannot be clearly mapped into motion behaviors, so that it is difficult for a robot to be aware of the potential damages.

The dependability of the system must be also understood. The ''passive'' safety is easy to understand: springs, rubber coverings, artificial skin, which have of course a real effect on reduction of damages, have also the additional property of being present in everyday life for the same purposes. The most obvious relation between cognitive ''fear'' related to physical injuries is the ''monster'' appearance of a robot, while the awareness of implemented safety systems could help. Consider ABS, ESP and the other systems for safety in the cars, which are probably the most ''familiar'' machines, and humanity came from a monster approach to the car and the train, present in lyrics and books of the first period of last century, to the human-like relation with this ''everyday'' ubiquitous machines. These safety devices are hidden and not clearly visible, but they add dependability and people know that they can trust on them: ''ubiquitous robotics'' as well as ubiquitous computing can be guaranteed if safety and dependability are guaranteed and understood.

The research in pHRI must consider any issue which could lead to define better evaluation criteria for the safety and dependability, considering ''scores'' or even ''cost functions'' to include the impact of different issues related to design and control of pHRI.

2.2. Safety standards and injury criteria

The main standard for robot safety in factories is the ANSI/RIA R15.06-1999 (American National Standard for Industrial Robots and Robot Systems – Safety Requirements). This standard addresses the requirements for personnel safety in industrial environments where robotic manipulators are employed. The complementary design standard ANSI/UL 1740 states hardware requirements and specifications, harmonized with R15.06: if the hardware is built in compliance with UL 1740, the safeguarding requirements in R15.06 are met too [22]. It must be pointed out, however, that there are no specifications for the case when robots and people have to share the operational space. Therefore, for unstructured anthropic domains this standard is not very useful.

The US Occupational Safety and Health Administration (OSHA) Manual discusses, in Section 4 of chapter 4, ''industrial robots and robot system safety''. The topic was addressed also in other countries: as an example, CAN/CSA Z434-03 is the Canadian standard on ''Industrial Robots and Robot Systems – General Safety Requirements''. The International Organization for Standardization (ISO) is currently revising the ISO 10218 standard (the international equivalent of R15.06), released in 1992, introducing new concepts in the world of industrial robot safety [23]. Because of the varying approach to standards writing across the world, it is essential to create a document that can be used in Europe, the Americas, and Asia. Among other things, these standards concern safety-related control circuitry performance for robots sharing their workspace with an operator, and involves also safety requirements for Autonomous Guided Vehicles (AGV).

Criteria for defining safety levels in human–robot interaction (inside and outside factories) are strictly related to the possible injuries caused by robots. Safety and dependability of the physical interaction should be evaluated considering all the different components of a robot, from mechanisms to actuators, from sensing to control. Note that recently some European robot manufacturers (ABB, KUKA, Reis Robotics) have included software modules that monitor through external sensing the Cartesian space around the robot and stop operations in case of danger.

For evident reasons, the automotive industry was the first to define quantitative measures, indices and criteria for evaluating injuries due to impacts. These sets of studies can be used as a starting point for safety evaluation in robotics. Most of the research about injury criteria is done in connexion with automobile crash testing where two distinct types of loading are observed concerning head injuries. The first type is a ''direct interaction'', i.e., a collision of the head with another solid object at appreciable velocity. This situation is generally characterised by large linear accelerations and small angular accelerations during the impact phase. The second type is an ''indirect interaction'', i.e., a sudden head motion without direct contact. The load is generally transmitted through the head-neck junction upon sudden changes in the motion of the torso and is associated with large angular accelerations of the head.

In order to quantify the injury severity produced by the impact with a robot, a scaling is needed. A definition of injury scaling developed by the automotive industry is the Abbreviated Injury Scale (AIS). If more than one part of the body is involved, the one with maximum injury severity is considered as the overall injury severity, which is indicated as Maximum AIS (MAIS). The type of injuries are divided in a classification

Table 1
AIS injury scale

AIS	Severity	Type of injury
0	None	None
1	Minor	Superficial
2	Moderate	Recoverable
3	Serious	Possibly recoverable
4	Severe	Not fully recoverable without care
5	Critical	Not fully recoverable with care
6	Maximum injury	Fatal

reported in Table 1: unfortunately, this table gives no indication how to measure injury. This is provided by so called severity indices.

Biomechanically motivated severity indices evaluated by impact tests are reviewed and discussed in [24]. Among the theoretical basements of these criteria, there are the Vienna Institute Index and the so-called Wayne State Tolerance Curve (WSTC). The latter builds upon data from fundamental experiments and represents the underlying reference for most of the proposed head-injury tolerance indices. The WSTC is a curve plotting head accelerations versus impact duration necessary to produce skull damages. Gadd [25] plotted the WSTC curve in log–log coordinates, obtaining a straight line of slope 2.5, and proposed accordingly a severity index (Gadd Severity Index) which is the integral of the head acceleration to the power of 2.5 during the relevant duration of collision.

For the head quite many criteria are available for the first type of loading (direct interaction). In the various interpretations of the WSTC, a model of the head is usually assumed in the form of a mass–spring–damper system.

The mostly used head severity index is the head injury criterion (HIC) [26]. The HIC focuses on the head acceleration and indicates that very intense acceleration is quite tolerable if very brief, while potentially harmful for pulse duration exceeding 10 or 15 ms (as time exposure to cranial pressure pulses increases, the tolerable intensity decreases).

The maximum power index (MPI) is instead the weighted change of kinetic energy of the human head before and after impacts, with the weighting carried out by two sensitivity coefficients in each direction. This injury criterion, which was developed for evaluation of protection helmets, is quite new and still a lot of research has to be carried out on it. However, it seems to be a promising alternative to the HIC since it is derived from the same underlying type of data (accelerations) but has physically relevant units.

In the maximum mean strain criterion (MSC), a mass–spring–damper model of the head is used and expanded by the presence of a second mass; the average length of the head is considered as a parameter too.

For torso injuries, the available criteria can be generally divided into four groups: acceleration based criteria, force based criteria, compression based criteria, and soft tissue based criteria. For neck injuries, frontal and rear impacts have different effects and thus they are addressed separately. In general, the mechanisms of injury of the human neck are potentially related to the forces and bending moments acting on the spinal column. Two severity indices for the neck were often considered: the neck injury criterion for frontal impacts and the neck injury criterion for rear impacts.

Unfortunately severity indices are not a direct scaling of injury but more of a limit between severe and non-severe injury. In order to cope with this deficit the National Highway Traffic Safety Administration specified empirical equations for converting e.g. HIC values to the probability of MAIS level. These conversions were carried out for several severity indices and provide a scaling of injury severity.

3. Mechanics and control issues for a safe pHRI

3.1. Issues for pHRI design

The discussed novel additional optimality criteria for pHRI lead to redesign robots starting from the mechanics. The intrinsic or passive safety cannot be underestimated: ''Rocks don't fly'' synthesizes the driving

force in the mechanical design of lightweight robots. The simple addition of a passive compliant covering in order to reduce impact loading is impractical and does not address the root cause of high impact loads due to the large effective inertia of most robotic manipulators. Finally, protective skins or helmets for humans are normal only in industrial domains, and not natural in anthropic domains. Operational tactics can also actively contribute to safety, by means of suitable control laws, and more sophisticated software architectures may overcome some limitations of mechanical structure. Indeed, control methods cannot fully compensate for a poor mechanical design, but they are relevant for performance improvement, reduced sensitivity to uncertainties, and better reliability. Modern actuation strategies, as well as force/impedance control schemes, seem to be anyway crucial in human–robot interaction. On the other hand, a more complete set of external sensory devices can be used to monitor task execution and reduce the risks of unexpected impacts. However, even the most robust architecture is endangered by system faults and human unpredictable behavior. This suggests to improve both passive and active safety for robots in anthropic domains.

3.2. Mechanics and actuation for pHRI

3.2.1. Intrinsic safety

The first important criterion to limit injuries due to collisions is to reduce the weight of the moving parts of the robot. A prototypical example of this is the design of the DLR-III Lightweight Robot [20], which is capable of operating a payload equal to its own weight (13.5 kg). Advanced light but stiff materials were used for the moving links, while motor transmission/reduction is based on harmonic drives, which display high reduction ratio and efficient power transmission capability. In addition, there is the possibility of relocating all the relevant weights (mostly, the motors), at the robot base, like it was done for the Barrett Whole Arm Manipulator (WAM) [27]. This is a very interesting cable-actuated robot, which is also backdrivable, i.e., by pushing on the links, it is possible to force motion of all mechanical transmission components, including the motors' rotors. In the case of a collision, the lighter links display lower inertia and thus lower energy is transferred during the impact. On the other hand, compliant transmissions tend to decouple mechanically the larger inertias of the motors from those of the links. The presence of compliant elements may thus be useful as a protection against unexpected contacts during pHRI. More in general, a lightweight design and/or the use of compliant transmissions introduce link [28] and, respectively, joint [29] elasticity.

In order to preserve performance while exploiting the potential offered by lightweight robot arms, one must consider the effects of structural link flexibility. Distributed link deformation in robot manipulators arises in the presence of very long and slender arm design (without special care on materials); notice that "link rigidity" is always an ideal assumption and may fail when increasing payload-to-weight ratio. On the other hand, in the presence of compliant transmissions, deformation can be assumed to be instead concentrated at the joints of the manipulator. Neglected joint elasticity or link flexibility limits static (steady-state error) or dynamic (vibrations, poor tracking) task performance.

Problems related to motion speed and control bandwidth must be also considered. Flexible modes of compliant systems prevent control bandwidths greater than a limit; in addition, attenuation/suppression of vibrations excited by disturbances can be difficult to achieve. Intuitively, compliant transmissions tends to respond slowly to torque inputs on the actuator and to oscillate around the goal position, so that it can be expected that the promptness of an elastically actuated arm is severely reduced if compliance is high enough to be effective on safety. From the control point of view, there is a basic difference between link and joint elasticity. In the first case, we have non-colocation between input commands and typical outputs to be controlled; for flexible joint robots, the co-location of input commands and structural flexibility suggests to treat this case separately.

3.2.2. Variable-impedance actuation

Very compliant transmissions may ensure safe interaction but be inefficient in transferring energy from actuators to the links for their fast motion. An approach to gain performance for guaranteed safety joint actuation is to allow the passive compliance of transmission to vary during the execution of tasks. The variable impedance approach (VIA) [30] is a mechanical/control co-design that allows varying rapidly and continuously during task execution the value of mechanical components such as stiffness, damping, and gear-ratio,

guaranteeing low levels of injury risk and minimizing negative effects on control performance. In this approach the best possible trade-off between safety and performance is desired. For a mechanism with given total inertia and actuator limits, one can formulate an optimal control problem to be used for comparing mechanical/actuation alternatives at their best control performance. One interesting formulation is the following: find the minimum time necessary to move between two given configurations (with associated motion and impedance profiles), such that an unexpected impact at any instant during motion produces an injury severity index below a given safety level. This is called the Safe Brachistochrone problem [31]. The optimal solution obtained analytically and numerically for single-dimensional systems shows that low stiffness is required at high speed and vice versa. This matches with intuition since most of the motion energy transfer from the motor should occur during the initial and final acceleration/deceleration phases. This ideal solution provides guidelines to be used also for multidimensional systems.

3.2.3. Distributed macro-mini actuation

Another approach to reduce manipulators arm inertia for safety, while preserving performance, is the methodology of distributed macro-mini actuation (DM^2) [21]. For each degree of freedom (joint), a pair of actuators are employed, connected in parallel and located in different parts on the manipulator. The first part of the DM^2 actuation approach is to divide the torque generation into separate low and high frequency actuators whose torque sum in parallel. Gravity and other large but slowly time-varying torques are generated by heavy low- frequency actuators located at the base of the manipulator. For the high-frequency torque actuation, small motors collocated at the joints are used, guaranteeing high-performance motion while not significantly increasing the combined impedance of the manipulator–actuator system. Finally, low impedance is achieved by using a series elastic actuator (SEA) [32]. It is important to notice that often high-frequency torques are almost exclusively used for disturbance rejection. The effective inertia of the overall manipulator is reduced by isolating the reflected inertia of the actuators while reducing the overall weight of the manipulator.

3.3. Control techniques for pHRI

Typically, current industrial robots are position-controlled. However, managing the interaction of a robot with the environment by adopting a purely motion control strategy turns out to be inadequate; in this case, a successful execution of an interaction task is obtained only if the task can be accurately planned. For unstructured anthropic domains, such a detailed description of the environment is very difficult, if not impossible, to obtain. As a result, pure motion control may cause the rise of undesired contact forces. On the other hand, force/impedance control [33] is important in pHRI because a compliant behavior of a manipulator leads to a more natural physical interaction and reduces the risks of damages in case of unwanted collisions. Similarly, the capability of sensing and controlling exchanged forces is relevant for cooperating tasks between humans and robots.

Interaction control strategies can be grouped in two categories; those performing indirect force control and those performing direct force control. The main difference between the two categories is that the former achieve force control indirectly via a motion control loop, while the latter offers the possibility of controlling the contact force to a desired value, thanks to the closure of a force feedback loop.

To the category of indirect force control belongs impedance control, where the position error is related to the contact force through a mechanical impedance of adjustable parameters. A robot manipulator under impedance control is described by an equivalent mass–spring–damper system, with the contact force as input (impedance may vary in the various task space directions, typically in a nonlinear and coupled way). The interaction between the robot and a human results then in a dynamic balance between these two ''systems''. This balance is influenced by the mutual weight of the human and the robot compliant features. In principle, it is possible to decrease the robot compliance so that it dominates in the pHRI and vice versa. Cognitive information could be used for dynamically setting the parameters of robot impedance, considering task-dependent safety issues.

Certain interaction tasks, however, do require the fulfilment of a precise value of the contact force. This would be possible, in theory, by tuning the active compliance control action and by selecting a proper reference location for the robot. If force measurements are available (typically through a robot wrist sensor), a

direct force control loop could be also designed. Note that, a possible way to measure contact forces occurring in any part of a serial robot manipulator is to provide the robot with joint torque sensors. The integration of joint torque control with high performance actuation and lightweight composite structure, like for the DLR-III lightweight robot [20], can help merging the competing requirements of safety and performance.

In all cases, the control design should prevent to introduce in the robot system more energy than strictly needed to complete the task. This rough requirement is related to the intuitive consideration that robots with large kinetic and potential energy are eventually more dangerous for a human in case of collision. An elegant mathematical concept satisfying this requirement is passivity. Passivity-based control laws [34], beside guaranteeing robust performance in the face of uncertainties, have thus promising features for a safe pHRI.

As already mentioned, compliant transmissions can negatively affect performance during normal robot operation in free space, in terms of increased oscillations and settling times. However, more advanced motion control laws can be designed which take joint elasticity of the robot into account. For example, assuming that the full robot state (position and velocity of the motors and links) is measurable, a nonlinear model-based feedback can be designed that mimics the result of the well-known "computed torque" method for rigid robots, i.e., imposing a decoupled and exactly linearized closed-loop dynamics [35]. Moreover, in robots with variable impedance actuation, the simultaneous and decoupled control of both the link motion and the joint stiffness is also possible in principle, reaching a trade-off between performance and safety requirements.

3.4. Real-time motion planning

Conventional robot motion planning is a typical off-line process that determines a feasible path (and a dynamically feasible timing), if one exists, connecting an initial and a final arbitrary robot configuration while avoiding obstacles. Complete knowledge of the geometry of the static environment is assumed. For high-dimensional configuration spaces (robots with many degrees of freedom) in crowded environments, the search for a feasible path is very complex and time-consuming; recently, probabilistic and randomized approaches have been developed to tackle this curse of dimensionality.

When dealing with trajectory (path + timing) planning in anthropic domains, the additional features of intelligibility and acceptability of robot motion should be considered. The planner should produce robot trajectories which are easily recognised by the user and "natural" for the task to be executed, in the sense that they are similar to the motion that would be performed by a human. In particular, they should not generate a sense of fear due to unexpected appearance, overly fast speeds/accelerations, and lack of visibility.

Within the context of off-line planning, one should include also an optimal or suboptimal definition of time variation for the stiffness properties of variable impedance actuation.

Indeed, the intrinsic nature of service robotics is to deal with unstructured, time-varying environments, for which a model is hardly available. The major part of dynamicity is typically due to the unpredictable motion of a human user. Therefore, the integration of a sensor-based on-line reactivity component into an off-line motion plan (needed for a global analysis of the scene) seems mandatory.

Sensors can be used to acquire local information about the relative position of a manipulator arm (or a navigating mobile robot) with respect to the human user (or with respect to other arms or robots, in which case proprioceptive sensing may be enough). Based on this, the planner should locally modify a nominal path so to achieve at least collision avoidance or, in more sophisticated strictly cooperating tasks, reach the correct rendezvous between the robot end-effector and, say, the human hand. Several reactive motion planning approaches exist in this context, mostly based on artificial potential fields [36] and their algorithmic or heuristic variations. One example is the elastic strip framework [37], where the local modification of motion is performed in a task-consistent manner, leaving primary (Cartesian) task execution possibly unaffected by obstacle avoidance. If the modification to the trajectories are given in the operational space, there is the need for an appropriate inverse kinematics system to give the reference values for the velocity/force controllers of the manipulator, possibly considering kinematic redundancy and/or dynamic issues. Another approach to reactive planning considers the on-line generation of the Cartesian path of multiple "control points" on the manipulator, possibly the closest one to obstacles [38].

In this framework, it is useful to encompass the entire physical structure of the robot and of the obstacles (human or not) with basic geometrical shapes (cubes, spheres etc.), which simplify distance computations.

The simplest modification of a nominal path in the proximity of an expected collision is to stop the robot. Even when a local correction is able to recover the original path, there is no guarantee, in general, that a purely reactive strategy may preserve task completion. For this, a global replanning based on the acquired sensory information may be needed.

4. Dependability in pHRI

4.1. Assessing dependability for pHRI

This section addresses the relevant dependability issues in pHRI, with emphasis on the robot side. Sensor capabilities and data fusion for inferring a correct characterization of the scene and of the people in the robot environment are discussed. Dependability of complex robotic systems in anthropic domains during normal operation is threatened by different kinds of potential failures or unmodeled aspects in sensors, control/actuation systems, and software architecture, which may result in undesirable behaviors. Due to the critical nature of pHRI, dependability must be enforced not only for each single component, but for the whole operational robot.

Dependability is an integrated concept that encompasses the following attributes [39]:

- Safety: absence of catastrophic consequences on the user(s) and the environment.
- Availability: readiness for correct service.
- Reliability: continuity of correct service, i.e., of completing tasks in a satisfactory manner.
- Integrity: absence of improper system alterations.
- Maintainability: ability to undergo modifications over time and repairs in case of failures.

There are, indeed, strict relationships among these concepts.

In all pHRI situations, safety of robot operation is essential, given the presence of humans in contact with or in the vicinity of the robots. In this context, safety can be rephrased as "absence of injury to humans in the robot's environment". Safety needs to be ensured both during nominal operation of the robot, as well in the presence of faults. In particular, it should be accepted that, in order to enforce a robot operation which is safe for the human, the completion of a programmed task may even be abandoned (this is also named survivability).

To be useful, a robot must also be always ready to carry out its intended tasks, and able to complete those tasks successfully. This is encapsulated in the dependability requirements of availability and reliability. There is evidently a trade-off between reliability/maintainability on one side, and safety on the other, since, in many applications, the safest robot would be one that never does anything. In some applications, however, the well being of humans requires robot availability and reliability. This is the case, for example, of robots used for surgery or for rescue operations.

Robot integrity is a pre-requisite for safety, reliability, and availability. Integrity relates to the robot's physical and logical resources, and requires appropriate mechanisms for protecting the robot against accidental and malicious physical damage, or corruption of its software and data.

Finally, availability cannot be achieved without due attention to maintainability aspects. These concern again both the physical and logical resources of the robot, which should be easy to repair and to upgrade.

4.2. Sensors and dependability

One major problem for the introduction of robots (in particular with mobile base) in unstructured environment is the possibility to rely on dependable sensors. Sensor data are needed for reactive planning, motion/force control, visual serving, fault diagnosis, and monitoring of safety levels. Due to the unstructured nature of anthropic domains and to the rather unpredictable movements of persons, a robot should be equipped with a complete set of sensors, including: range, proximity, touch, vision, sound, temperature, and so on. The selection, arrangement, and number of sensors (as well as their single reliability) contribute to the measure of dependability of a manipulator for interaction tasks.

Intelligent environments can be considered as dual to a robot equipped with multiple sensors. On one hand, a properly sensorized environment incorporates the tools to perceive and understand what is happening, providing useful data to people and robots present in the environment. In this way, many robots can share the same sensory system. However, the large work needed for setting up such a global sensory system makes this approach unpractical for an anthropic environment. On the other hand, robots with built-in sensors provide a versatile platform, since their operation is not limited to a specific area.

The construction of a good model of humans interacting with the robot is certainly one of the main purposes of a sensory system for pHRI: vision and other proximity sensors must be able to map the position of critical actors in the scene. These sensors must be robust to changing of environmental conditions like lighting, dust and other sources of uncertainty.

Sensor-based modeling of a person was addressed in the literature [40], in order to obtain enough information about human intentions, by tracking and interpreting movements and gestures. For this, vision is the main sensory modality and stereo cameras can provide also depth information. Basic capabilities include locating people in the operational space, predicting their incipient motion (based on a kinematic/dynamic model of the human) and detecting faces. Anticipated perception [41] and anticipatory movements can help the robot to keep adequate postures during the interaction with human beings. An articulated human model [42] should be considered, using at least simple skeletons of body components (neck, shoulders, torso, arms, forearms, hip, legs).

Active exploration and selectivity of features are two important characteristics of the human vision which should be replicated in robot vision. With an on-board camera, robot motion can be used for active acquisition of data on human–robot interaction. As a minimum, a localization/detection module should evaluate direction and distance of people present in the workspace. "Safety volumes" [38] can be generated so as to constrain the behavior of the robot, in terms of speed and compliance, and depending on the specific phase of the interaction. In particular, this is relevant in the case of mobile service robots approaching humans (e.g., encountering them in a corridor). The selective visual tracking of the head is crucial [43], considering the presented injury criteria, but also monitoring arm positions is important, due to the possibility of dangerous collisions during cooperative tasks. A more precise localization system can be built using radio frequency identification tags or markers.

For closer human–robot interaction, such as supporting people, handing over objects or shaking hands-all "physical" events, force/torque sensors and small "tactile" sensors distributed along the structure may be needed (just like in dextrous robotic hands). Of course, it is necessary to distinguish between intentional contacts and accidental collisions with human body and head.

Fusion of the information coming from multiple sensors may help in providing a coherent and reliable description of the world surrounding the robot. In general, it is required to integrate sensor information based only on approximate models of the environment. Data fusion is particularly important when monitoring contacts, e.g., for selecting impedance parameters or for determining the most dangerous "control points" on the robot to be driven away from a human with higher priority. Unfortunately, there has been little work on achieving the fusion of contact and visual information.

Taking as an example the problem of simultaneous localization and mapping (SLAM) [44] for a mobile service robot, the input from different sensors (laser scan, ultrasonic, vision) may be combined in order to drive prediction algorithms. The latter are based on Bayesian estimation, Kalman filtering, particle filters and so on, and cope with the occasional and systematic uncertainties of the observations.

Artificial intelligence (AI) techniques can also be adopted for data fusion, e.g., fuzzy sets, neural networks, or first order logics. For the use of AI systems, there is often the need to define "symbolic" quantities that capture some relevant state of the human–robot interaction and/or of the environment to be used for high-level reasoning. Finally, inference systems may organize sensory sources and data taking into account the information about the specific phases of an interaction task. Indeed, only AI methods complying with real-time constraints are of interest for pHRI.

4.3. Fault handling and dependability

The possibility of conferring a proper degree of autonomy and safety to robots strongly depends on the capability to properly manage the possible occurrence of unexpected events, as failures or abrupt changes

of the environment. To preserve the safety of humans interacting with robots during the execution of inter-action tasks, fault handling and fault tolerant control have to be considered as fundamental functionalities [45]. Dependability is related on the ability of the system to cope with failures.

As an example, a model of failure taxonomy has been presented in [46] with the reference to field robots. Reliability of the system is evaluated in terms of physical failures (of the effectors, sensors, power source, control system, and so on) and human communication failures. It must be pointed out, however, that for application domains with physical human–robot interaction, the picture is even more complex.

To ensure acceptable levels of robot dependability attributes in pHRI, it is useful to define explicitly the types of faults that can affect the robot, and that need to be taken into account during development and deployment. These can be very broadly described in terms of three non-disjoint fault classes:

- physical (or internal) faults, including both natural hardware faults and physical effects due to the environment (damage of mechanical parts, actuators and/or sensors faults, power supply failures, control unit hardware/software faults, radiation, electromagnetic interference, heat, etc.);
- interaction (or external) faults, including issues related to human-to-robot and robot-to-robot cooperation, robustness issues with respect to operation in an open, unstructured environment (such as sudden environmental changes and disturbances not usually acting during the normal system operation or exceeding their normal limits), and malicious interference with the robot's operation;
- development faults, which may be introduced, usually accidentally, during the design or implementation of the hardware and software components of the robot.

All three faults classes need to be considered, with more or less emphasis depending on the application. One particularly delicate aspect in the context of robotics is that of development faults affecting the domain-specific knowledge embodied in robot's world models and the heuristics in decisional mechanisms.

Achieving dependability requires the application of a sequence of activities for dealing with faults [39]:

- fault prevention, which aims to prevent the occurrence or introduction of faults;
- fault removal, which aims to reduce the number and severity of faults;
- fault detection and isolation, which aims at recognizing the occurrence of a fault and characterizing its type;
- fault tolerance, which aims to avoid service failures in the presence of faults;
- fault forecasting, which aims to estimate the present number, the future incidence, and the likely consequences of faults.

Fault prevention and fault removal are collectively referred to as fault avoidance, i.e., how to build a system that is fault-free. With respect to development faults, fault avoidance requires a rigorous approach to design through proven software engineering principles and the application of verification and validation procedures, such as formal verification and testing.

Fault detection and isolation are parts of fault diagnosis. An early diagnosis allows an appropriate response of the system and prevents the propagation of the fault effects to critical components in the system. The robotic system has to be monitored during its normal working conditions so as to detect the occurrence of failures (fault detection), recognize their location and type (fault isolation), as well as their time evolution (fault identification). Fault diagnosis methodologies are based on hardware redundancy, in the case of duplicating sensors, or on analytic redundancy, in the case that functional relationships between the variables of the system (usually obtained from the available mathematical model) are exploited. Usually, the output of a fault diagnosis algorithm is a set of variables sensitive to the occurrence of a failure (residuals), affected by a signature in the presence of a fault (fault signature). Then, the information from the signatures is processed to identify the magnitude and the location of the fault. Sometimes it is also possible to achieve a one-to-one relation between faults and residuals (decoupling), so that fault isolation is obtained, without further processing, as a byproduct of detection. Existing analytical fault diagnosis techniques include observer-based approaches, parameter estimation techniques, and algorithms based on adaptive learning techniques or on soft-computing methodologies.

In practice, avoiding all possible faults is never fully achievable. Fault tolerance and fault forecasting are collectively referred to as fault acceptance, i.e., how to live with the fact that the robotic system may contain or

is actually affected by faults. In pHRI, fault acceptance requires tolerance (or robustness) with respect to adverse environmental situations and other interaction faults, and incorporation of redundancy to tolerate faults affecting robotic hardware or software. The concept of redundancy may be cast into a modular design philosophy, both hardware and software, that may guarantee that the effects of local faults remain internal to the modules, and also permits the reconfiguration of the system.

In particular, fault tolerant control strategies can be separated into passive and active methods (possibly, to be combined). The passive approaches are based on the adoption of robust control techniques to ensure that the controlled system remains insensitive to certain fault categories, considered as modelling errors and disturbances. In the active approaches, when a failure occurs and is diagnosed (the fault has been isolated and possibly identified), the controller is reconfigured in order to preserve some properties of the controlled system, even though with degraded performance (adaptive control approaches belong to this class).

Another important aspect in the development of fault tolerant systems relies on the adoption of critical components redundancy. For robotic systems, redundancy can be introduced by adopting additional actuators, as in the case of duplicating joint actuators in spatial robots, or multiple sensory devices. Additionally, one may exploit kinematic redundancy of a manipulator; in such a case, a failed joint can be braked and the task accomplished by suitably modifying the trajectories of the healthy joints.

In the case of robotic systems interacting with humans, an intrinsically safe interaction and high tolerance to unexpected collisions can be guaranteed by imposing a suitable programmable compliant behavior of the robotic system, e.g., via impedance control strategies. When a failure occurs, the robotic system should reach a configuration maximally safe for the humans.

To judge whether the adopted robustness and tolerance techniques are necessary and sufficient, the achieved dependability needs to be assessed by an appropriate combination of analysis (e.g., failure modes, effects, and criticality analysis (FMECA), hazard analysis) and evaluation (e.g., through stochastic modeling or experimental fault injection).

4.4. Control architecture and dependability

Dependability of the robot control software for pHRI calls for a modular and hierarchical architecture, which is also advantageous for testing the single components and isolating possible faults so as to achieve operating robustness (availability, reliability, and maintainability [8]).

Due to the need for continuous monitoring the environment and robot operation, as well as for on-line changes in planning Cartesian (and stiffness) trajectories, the operating system of a control architecture for pHRI must run in "real-time". Since tasks need to be executed within a specified time constraint, multitasking, prioritized process threads, and multiple interrupt levels are to be considered.

Deploying a control architecture for pHRI requires the proper organisation of the various software components. The type of onboard software for such robots will be very diverse, from low-level control loops up to high-level planners. Nonetheless, the expected overall properties of such an architecture are:

- Programmability: the robot should be able to achieve multiple tasks described at an abstract level. Its basic functionalities should therefore be easily combined according to the task to be executed.
- Autonomy and adaptability: the robot should be able to carry out its actions and to refine or modify the task and its own behavior according to the current goal and execution context as it perceives it.
- Reactivity: the robot has to take into account events with time bounds compatible with the correct and efficient achievement of its goals (including its own safety) and the dynamics of the environment.
- Consistent behavior: the reactions of the robot to events must be guided by the objectives of its task.
- Robustness: the control architecture should be able to cope with failures, in particular by exploiting the redundancy of the processing functions and subsystems. Robustness will require the control to be decentralized to some extent.

The above requirements imply the coexistence of both deliberative and reactive behaviors in the system. The architecture should therefore embed interacting subsystems that perform according to different temporal properties. In general, the implementation of several task-oriented and event-oriented closed-loops for

achieving both anticipation capacities and real-time behavior cannot be done on a single system with homogeneous processes, due to their different computational requirements.

A global architecture is needed, that enables the integration of processes with different temporal properties and different representations. One example is given by the LAAS architecture [47], which is composed by two hierarchical levels.

The lowest level, namely the functional level, contains all the basic built-in functionalities of the system. Processing functions and control loops (e.g., image processing, obstacle avoidance, motion control, etc.) are encapsulated into controllable communicating modules. To make this level as hardware-independent as possible, and hence portable from one robot to another, it is interfaced with the sensors and effectors through a logical robot level. The modules are activated by the next level according to the task.

The highest level is the decision layer, composed of a supervisor and a planner, which constitutes the main software component for autonomy. This level includes the capabilities for producing the task plan and supervising its execution, while being at the same time reactive to events from the functional level. The coexistence of these two features, a time-consuming planning process and a time bounded reactive execution process, poses the key problem of their interaction and their integration to balance deliberation and reaction at the decisional level.

Moreover, a solution to guarantee the safety properties is to integrate in the architecture a module that formally checks the validity of the low-level commands sent to the physical system and prevents the robot from entering an unsafe state. This check must verify consistency during execution without affecting the system basic functionalities: this mechanism is part of an execution control layer. Despite its clear role in the architecture, the execution control layer can actually be embedded in the functional level, given its intricate interaction with its components. This results in an overall two-level architecture.

5. Benchmarks for pHRI

5.1. Examples of case studies

The definition of benchmarks is a general problem which becomes even more crucial in service robotics [48]. For instance, a domestic living room was set up at the 2002 IEEE-RSJ IROS Conference in Lausanne for a contest among robotic vacuum cleaners. Similarly, in the RoboCup competition, rescue robots are compared in predefined emergency arenas of increasing complexity.

Being anthropic domains quite different, it is difficult at this stage to propose a unique environment to assess dependability and performance of robots. Therefore, we present here some simple case studies highlighting how safety issues are taken into account in practice.

Incidentally, we note that some service robots either address the safety issues with simple solutions (bumpers in robotic vacuum cleaners) or present small intrinsic risks (such as small entertainment or pet robots).

An example of physically interacting robot providing power augmentation to humans workers is the Cobot [49]. In one of its basic implementations, it is a wheeled robotic platform that supports (typically heavy) parts to be manipulated by an operator. Virtual guiding surfaces are created, directing the constrained motion toward the appropriate environment location. The virtual guiding surfaces can be programmed in space and time and blended one into another. An assembly assist tool is made up of a guidance unit (the cobot) as well as conventional task-dependent tooling (e.g., a door loader). Ergonomics is the performance criterion, with an improved inertia management leading to smaller operator applied forces. Safety is addressed via the intrinsic passivity of the cobot: the maximum energy in the system is limited by the human's capability. Also cobots with power assist have been developed: although these robots are not fully passive, safety is still preserved by appropriately limiting the power of the assisting motor. In this case study, the safety problem was solved enforcing a "human-in-the-loop" strategy.

Another example of application where safety is considered as a primary task are exoskeletons [5].

Related to dependability and robustness of safe robots, the possible failure modes of a simple robot with a Variable Impedance Actuation, based on antagonistic arrangement on nonlinear elastic elements, have been analysed in [50] under possible failures of some of its components. The ability of the system to remain safe in spite of failures has been compared with that of other possible safe-oriented actuation structures, namely, the

SEA and to the DM^2 actuation scheme. In order to obtain a meaningful comparison, optimized SEA and DM^2 implementations have been considered, with equal rotor and link inertias, yielding the same minimum-time motion performance for the considered task. Under the same failure modes, both SEA and DM^2 lead to higher HIC values. An explanation of the apparently superior fail-safety characteristics of the antagonistic VIA is that such scheme achieves comparable nominal performance by employing two motors each of much smaller size than what necessary in the SEA and DM^2. The basic stiff-and-slow/fast-and-soft idea of the VIA approach seems therefore to be more effective for realistic models of antagonistic actuation.

Notice again that the problem of collisions is a central topic for research and experiment in pHRI, both for collision avoidance and for robot reconfiguration after collisions. Related to the second case, collision detection in the absence of external sensing devices can be realized in different ways by suitably comparing commanded motor torques and measured proprioceptive signals [51,52]. A particularly efficient algorithm that uses only encoder positions is based on the monitoring of the generalized momentum of the mechanical system [53],[54], which allows also to identify (isolate) the colliding link on the robot.

Once the collision has been detected (more or less as a system fault), the robot may simply be stopped by braking or applying high-gain position feedback on the current joint position. However, the robot will remain in the vicinity of the collision zone with the human, producing thus a sensation of permanent danger. In [55], a different strategy has been implemented on a lightweight robot arm, by determining a direction of safe post-impact motion for the robot from the same signal used for collision detection.

We note finally that, if the collision is assumed to occur at the end-effector level (say, between the robot tool and the human user) kinematic redundancy of the arm may be used to minimize the instantaneous effect of an impact [56]. In fact, while executing a desired end-effector trajectory, the arm may continuously change its internal kinematic configuration in order to minimize the inertia seen at the end-effector.

5.2. Revising safety measures

Based on the previous discussion and on considered applications, it is clear that an assessment of the safety level for physical human–robot collisions is mandatory. Revised safety metrics, including the effects of the adopted reaction strategies on the reduction of risk, should be considered [57,58]. The common injury indices reviewed in Section 2 originate in fact from fields other than robotics (e.g. automotive crash test, development of protection helmets) and were developed and tailored for those specific applications. Although they constitute a useful basis for starting the development and evaluation of safe robotic concepts, special attention has to be paid to the question whether the conditions under which these indices were formulated are indeed satisfied (and are general enough) in their robotic application.

As an example, consider the case of the head injury criterion (HIC). This widely adopted index in automotive crash tests was also successfully used so far in the pioneering pHRI literature in order to evaluate various hardware concepts for inherent safety mechanical design (VIA and DM^2). This index is based on the evaluation of the time profile of the head acceleration. In the case of a car crash, the main source of injury for the head is the acceleration transmitted from the vehicle, through the body of the person, to its head. It is very likely that no direct impact of the head occurs during an accident and consequently the acceleration is the only quantity which can always be reliably measured (e.g. by a test dummy). Nevertheless, the index applies also for impacts, if the head can freely move (accelerate) after the impact. In robotics however, the most dangerous injury would probably be produced always by a direct impact of the robot with the human's head. Moreover, it is possible that the head (or another part of the body) is even squeezed between robot and environment or between parts of the robot, leading to only reduced accelerations (and hence to an uncritical value of the HIC), but nevertheless constituting a high injury risk. This observation suggests that an index related, e.g., to the impact force or to the transmitted energy could describe more precisely the various types of injury which can arise in pHRI.

This punctual example is given as a motivation for the need of further research for establishing a set of indices and measures that quantify as completely as possible the potential danger emanating from a robot. One obvious possibility is to evaluate various existing indices (and possibly proposing new ones) based on simulation studies using existing models of the human body and of test robotic scenarios. Of course, the validation of the criteria is very problematic, one still has to rely here on empirical data (relating forces, accelerations,

impact energy to real human injury) available from other fields. It is the implicit goal of the safe-robotics research activities, to prevent the emergence of such an database of eerily accidents in robotics. The availability of standardized, well founded, and reliable criteria will considerably pave the way for making robot–human interaction safe and hence feasible and attractive from a practical point of view.

An example of extension of the indices to the peculiar problem of pHRI is the Manipulator Safety Index [59], derived from a combination of the well-known HIC index with a simplified analytical robot model-based on manipulator dynamics. The manipulator's effective inertia, interface stiffness, and possible impact velocity result in different values of such a criterion in different poses and directions.

Although some HRI taxonomies were proposed (see, e.g., [60]), the main issues considered (autonomy, human/robot presence ratio, level of shared interaction, available sensors and their fusion, task criticality) do not include safety in physical interaction. A taxonomy for the evaluation of pHRI, with emphasis on safety and dependability issues is still missing. According to the discussion in this document, one may provide ''scores'' to the safety/dependability of the following robot components and functionalities: lightweight mechanical design, passive soft covering, (variable) compliant actuation/transmission, complete sensor suite/fusion, human-oriented off-line planning, reactive on-line planning, stable force/impedance control, motion control performance, fault diagnosis and tolerance, collision avoidance or detection, collision reaction tactics, modular control architecture. These individual scores should be combined with suitable weights and evaluated on a large sets of consistent experiments. A checklist associated to a typical robotic task involving pHRI should be considered. The scenario may include a robot manipulator mounted on a mobile base, used as an assistant for collecting an object in a large environment and handing it over to a human without damages or harm to humans present in the workspace.

Acknowledgement

References

[1] O. Khatib, K. Yokoi, O. Brock, K. Chang, A. Casal, Robots in human environments: basic autonomous capabilities, International Journal of Robotics Research 18 (7) (1999) 684–696.
[2] http://www.wtec.org/robotics/us_workshop/June22/MilitaryRobots%20WTEC.pdf.
[3] Proceedings of the Italy–Japan 2005 Workshop The Man and the Robot: Italian and Japanese approaches, http://www.robocasa.net/workshop2005/index.php?page=program.
[4] E. Rogers, Introduction to Human–Computer Interaction (HCI), RAS/IFRR Summer School on ''Human–Robot Interaction, http://www.cas.kth.se/ras-ifrr-ss04/material/rogers-hci-intro.pdf.
[5] http://people.na.infn.it/~ethicbots/index.htm.
[6] http://www.neurobotics.info/.
[7] http://bleex.me.berkeley.edu/bleex.htm.
[8] http: //othello.mech.northwestern.edu/~peshkin/cobot/.
[9] H. Hashimoto, Intelligent interactive space – integration of IT and robotics, in: IEEE Workshop on Advanced Robotics and its Social Impacts, 2005.
[10] N. Reeves, C. Nass, The Media Equation: How People Treat Computers, Television and New Media Like Real People and Places, Cambridge University Press, Cambridge, Mass., 1996.
[11] K. Severinson-Eklundh, A. Green, H. Huettenrauch, Social and collaborative aspects of interaction with a service robot, in: Workshop on ''Robot as Partner: an exploration of social robot'', IEEE/RSJ International Conference on Intelligent Robots and Systems, Lausanne, Switzerland, 2002.
[12] I. Calvino, Six Memos for The Next Millennium, Harvard University Press.
[13] J.L. Burke, R. Murphy, E. Rogers, V.L. Lumelsky, J. Scholtz, Final report of the DARPA/NSF Interdisciplinary Study on Human–Robot Interaction.
[14] ''Safely among us'', IEEE R&A Magazine, vol. 11, no. 2, Special Issue on Dependability of Human-Friendly Robots, June 2004.
[15] H. Huettenrauch, K.S. Eklundh, Fetch-and-carry with CERO: observations from a long-term user study with a service robot, in: 11th IEEE International Workshop on Robot and Human Interactive Communication, 2002.
[16] http://www.cogniron.org/.

[17] http://robocare.istc.cnr.it/.

[18] Technical challenge for dependable robots in human environments, in: G. Giralt, P. Corke (Eds.), Proceedings of IARP/IEEE-RAS Joint Workshop, Seoul, Korea, 2001.

[19] A. Albu-Schäffer, A. Bicchi, G. Boccadamo, R. Chatila, A. De Luca, A. De Santis, G. Giralt, G. Hirzinger, V. Lippiello, R. Mattone, R. Schiavi, B. Siciliano, G. Tonietti, L. Villani, Physical human–robot interaction in anthropic domains: safety and dependability, in: 4th IARP/IEEE-EURON Workshop on Technical Challenges for Dependable Robots in Human Environments, Nagoya, J, July 2005.

[20] G. Hirzinger, A. Albu-Schaeffer, M. Hahnle, I. Schaefer, N. Sporer, On a new generation of torque controlled lightweight robots, in: IEEE International Conference of Robotics and Automation, Seoul, K, 2001.

[21] M. Zinn, O. Khatib, B. Roth, J.K. Salisbury, A new actuation approach for human-friendly robot design, in: International Symposium on Experimental Robotics, S. Angelo d'Ischia, I, July 2002.

[22] http://www.roboticsonline.com/public/articles/archivedetails.cfm?id=497.

[23] http://www.roboticsonline.com/public/articles/archivedetails.cfm?id=1215.

[24] S. Haddadin, Bewertungskriterien und Regelungsstrukturen fuer eine sichere Mensch-Roboter Interaktion, Diplomarbeit, DLR, Oberpfaffenhofen, D, 2006.

[25] C.W. Gadd, Use of weighted impulse criterion for estimating injury hazard, in: 10th Stapp Car Crash Conference, New York, 1966.

[26] J. Versace, A review of the severity index, in: Proceedings of the 15th Stapp CarCrash Conference, New York, 1971.

[27] http://www.barretttechnology.com/robot/products/arm/armfram.htm.

[28] A. De Luca, Feedforward/feedback laws for the control of flexible robots, in: IEEE International Conference of Robotics and Automation, San Francisco, CA, USA, 2000.

[29] A. De Luca, P. Lucibello, A general algorithm for dynamic feedback linearization of robots with elastic joints, in: IEEE International Conference on Robotics and Automation, Leuven, B, 1998.

[30] A. Bicchi, G. Tonietti Fast and 'soft-arm' tactics, IEEE R&A Magazine, vol. 11, no. 2, Special Issue on Dependability of Human-Friendly Robots, 2004.

[31] A. Bicchi, S. Lodi Rizzini, G. Tonietti, Compliant design for intrinsic safety: general issues and preliminary design, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, Maui, HI, USA, 2001.

[32] G. Pratt, M. Williamson, Series elastic actuators, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, Pittsburgh, PA, USA, 1995.

[33] B. Siciliano, L. Villani, Robot Force Control, Kluwer Academic Publishers, Boston, MA, 1999.

[34] A. Albu-Schäffer, C. Ott, G. Hirzinger, A unified passivity-based control framework for position, torque and impedance control of flexible joint robots, in: 12th International Symposium of Robotics Research, San Francisco, CA, USA, 2005.

[35] M.W. Spong, Modeling and control of elastic joint robots, ASME Journal of Dynamic Systems, Measurement, and Control 109 (1987) 310–319.

[36] O. Khatib, Real-time obstacle avoidance for robot manipulators and mobile robots, International Journal of Robotics Research 5 (1) (1986) 90–98.

[37] O. Brock, O. Khatib, Elastic strips: a framework for motion generation in human environments, International Journal of Robotics Research 21 (12) (2002) 1031–1052.

[38] A. De Santis, P. Pierro, B. Siciliano, The virtual end-effectors approach for human–robot interaction, in: 10th International Symposium on Advances in Robot Kinematics, Ljubljana, SLO, 2006.

[39] A. Avizienis, J.-C. Laprie, B. Randell, C. Landwehr, Basic concepts and taxonomy of dependable and secure computing, IEEE Transactions on Dependable and Secure Computing 1 (1) (2004) 11–33.

[40] M. Frigola, A. Rodriguez, J. Amat, A. Casals, Robust modeling of a body for human–robot interaction, in: IEEE International Conference on Robotics and Automation, Workshop on "The Social Mechanisms of Robot Programming by Demonstration", Barcelona, E, 2005.

[41] E. Datteri, G. Teti, C. Laschi, G. Tamburrini, P. Dario, E. Guglielmelli, Expected perception: an anticipation-based perception-action scheme in robots, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas, Nevada, USA, 2003.

[42] K.M. Cheung, S. Baker, J.K. Hodgins, T. Kanade, Markerless human motion transfer, in: 2nd International Symposium on 3D Data Processing, Visualization and Transmission, September 2004.

[43] P. Viola, M. Jones, Robust real-time face detection, International Journal of Computer Vision 57 (2) (2004) 137–154.

[44] J.J. Leonard, H.F. Durrant-Whyte, Simultaneous map building and localization for an autonomous mobile robot, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, Osaka, Japan, 1991.

[45] F. Caccavale, L. Villani, Fault Diagnosis and Fault Tolerance for Mechatronic Systems: Recent Advances, Springer Tracts in Advanced Robotics, vol. 1. Heidelberg, Germany, 2003.

[46] J. Carlson, R.R. Murphy, A. Nelson, Follow-up analysis of mobile robot failures, in: IEEE International Conference of Robotics and Automation, New Orleans, LA, USA, 2004.

[47] R. Alami, R. Chatila, S. Fleury, M. Ghallab, F. Ingrand, An architecture for autonomy, International Journal of Robotics Research (1998).

[48] http://www.robot.uji.es/benchmarks/.

[49] R.B. Gillespie, J.E. Colgate, M.A. Peshkin, A general framework for cobot control, IEEE Transactions on Robotics and Automation 17 (4) (2001).

[50] G. Boccadamo, R. Schiavi, S. Sen, G. Tonietti, A. Bicchi, Optimization and fail-safety analysis of antagonistic actuation for pHRI, in: 1st EUROS Symposium, Palermo, I, 2006.

[51] K. Suita, Y. Yamada, N. Tsuchida, K. Imai, H. Ikeda, N. Sugimoto, A failure-to-safety 'kyozon' system with simple contact detection and stop capabilities for safe human-autonomous robot coexistence, in: IEEE International Conference on Robotics and Automation, Nagoya, Japan, 1995.
[52] Y. Yamada, Y. Hirasawa, S. Huang, Y. Uematsu, K. Suita, Human–robot contact in the safeguarding space, IEEE/ASME Transactions on Mechatronics 2 (4) (1997) 230–236.
[53] A. De Luca, R. Mattone, Sensorless robot collision detection and hybrid force/motion control, in: IEEE International Conference on Robotics and Automation, Barcelona, E, 2005.
[54] H.-B. Kuntze, Ch.W. Frey, K. Giesen, G. Milighetti, Fault tolerant supervisory control of human interactive robots, in: IFAC Workshop on Advanced Control and Diagnosis, Duisburg, D, 2003.
[55] A. De Luca, A. Albu-Schaeffer, S. Haddadin, G. Hirzinger, Collision detection and safe reaction with the DLR-III lightweight manipulator arm, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, Beijing, PRC, 2006.
[56] I.D. Walker, Impact configurations and measures for kinematically redundant and multiple armed robot systems, IEEE Transactions on Robotics and Automation 10 (5) (1994).
[57] K. Ikuta, H. Ishii, M. Nokata, Safety evaluation method of design and control for human-care robots, International Journal of Robotics Research 22 (7/8) (2003) 281–297.
[58] J. Heinzmann, A. Zelinsky, Quantitative safety guarantees for physical human–robot interaction, International Journal of Robotics Research 22 (7/8) (2003) 479–504.
[59] M. Zinn, O. Khatib, B. Roth, Human–robot physical interaction: design for safety and performance, in: Lecture by O. Khatib at the IEEE-RAS/IFRR International School of Robotics Science on "Robot Design", Tokyo, J, 2005.
[60] H. Yanco, J.L. Drury, A taxonomy for human–robot interaction, in: AAAI Fall Symposium on Human–Robot Interaction, Falmouth, Massachusetts, November 2002.
[61] http://www.piaggio.ccii.unipi.it/phridom/index.htm.
[62] http://www.phriends.eu.