

Safety Issues in Human-Robot Interactions

Milos Vasic¹ and Aude Billard²

Abstract—Safety is an important consideration in human-robot interactions (HRI). Robots can perform powerful movements that can cause hazards to humans surrounding them. To prevent accidents, it is important to identify sources of potential harm, to determine which of the persons in the robot’s vicinity may be in greatest peril and to assess the type of injuries the robot may cause to this person. This survey starts with a review of the safety issues in industrial settings, where robots manipulate dangerous tools and move with extreme rapidity and force. We then move to covering issues related to the growing numbers of autonomous mobile robots that operate in crowded (human-inhabited) environments. We discuss the potential benefits of fully autonomous cars on safety on roads and for pedestrians. Lastly, we cover safety issues related to assistive robots.

I. INTRODUCTION

Robots can produce powerful and very rapid movements through a large operational space. Hazard threats arise from unintended contact between these robots and humans. The forthcoming paths of robots or robots’ arms are difficult to predict (e.g., due to changing operational requirements). Operators can be required to work in close proximity to the robot system, while the machine actuators are being powered. Furthermore, the operating spaces of two or more robots can overlap, representing a threat to a human worker from multiple sources.

Nowadays, humans work in close cooperation with robots more than ever. In the resulting situations, contact is unavoidable between humans and robots. Indeed, in many instances, it is actually desired. However, every contact creates the potential for an accident. Careful thought needs to be given to hazard assessment. Once the hazards are known, they can be eliminated or reduced by design, safeguarding, control and other methods.

Industrial robot interactions between humans and robots are perhaps the most well-studied in current literature. However, the emerging field of mobile and autonomous robots brings new threats that have not yet been equally well-assessed.

When it comes to human safety, accident prevention can always be improved. There have been many accidents caused by HRI in the past. In the years to come, the number of robots surrounding humans will greatly increase. Hence, the potential danger from robots is greater than ever. The purpose of this survey is to review possible hazards and overview the methods used for accident risk reduction, as well as point out

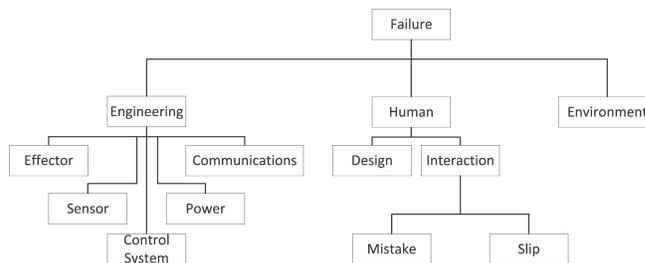


Fig. 1. Taxonomy of failures, adapted from [2]. Failures can be caused by engineering or human worker errors, as well as environmental conditions.

the risks that did not exist before, but are being introduced with the development of modern robotic systems.

II. HAZARDS OVERVIEW

In order to improve safety in HRI, one needs to determine a) where the biggest danger lies, b) who is the most endangered person in the interaction with a robot, c) what are the consequences of potential injuries, and d) which factors have the greatest impact on safety? In this section we address each of these questions in turn.

A. Sources of injuries

According to [1], the causes of accidents caused by robots can be grouped into three main categories: engineering errors, human mistakes and poor environmental conditions. Engineering errors include errors in the robot’s mechanics (loose connections across parts, faulty electronics), errors made by the controller (programming bugs, faulty algorithm), etc. As a consequence, robots might, for example, fail to stop, or a robot arm might achieve high, uncontrolled speed, abrupt motion or acceleration. Accident caused by these errors cannot be predicted even by the most attentive human operator. On the other hand, human accidents, which are more controllable, happen due to various factors, such as inattention, fatigue, inobservance of the guarding procedures, inadequate training programs or incorrect procedures for initial robot start-up. Adverse environmental factors refer to extreme temperature, poor sensing in difficult weather or lighting conditions, all of which can lead to incorrect response by the robot. Fig. 1 depicts the classifications of accident sources mentioned above.

B. Endangered personnel

The person who operates the robot is the most at risk. A comprehensive report based on cause-effect analysis of 32 accidents is presented in [3]. Its results show that robot operators were subject to injuries in 72% of the reported accidents. Maintenance workers accounted for 19%

¹M. Vasic is with the Distributed Intelligent Systems and Algorithms Laboratory, School of Architecture, Civil and Environmental Engineering, Ecole Polytechnique Federale de Lausanne, 1015 Lausanne, Switzerland milos.vasic@epfl.ch

²A. Billard is with the Learning Algorithms and Systems Laboratory, School of Engineering, Ecole Polytechnique Federale de Lausanne, 1015 Lausanne, Switzerland aude.billard@epfl.ch

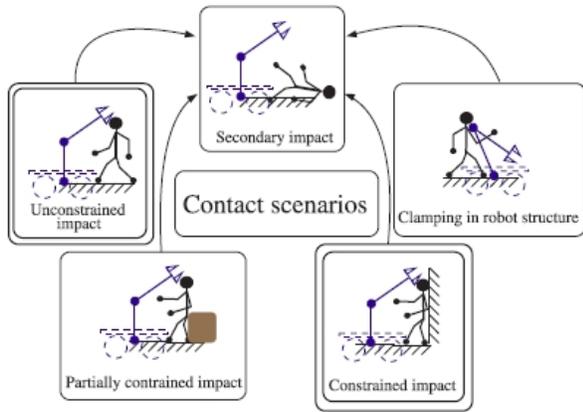


Fig. 2. Classification of undesired contact scenarios between human and robot [4].

of accidents, whereas programmers were least prone to accidents (9% of cases). This directly corresponds to the amount of time a person spends in the proximity of a robot, as well as his or her level of expertise. Robot operators are usually instructed to operate the robots in routine-like situations only, and are therefore unprepared for unexpected robot behavior. Maintenance workers are usually better trained to handle such uncertainty. In many situations, however, they are called when it is already known that the robot is not functioning properly. They therefore maintain a higher level of attention and are more precautious. Nevertheless, many maintenance workers get injured. These injuries are generally due to human mistakes, such as when another maintenance worker activates a robot system to test it, while the first worker is still in the robot cell. The robot programmers have exceptional knowledge of robot operation, so their injuries often fall into two categories. Most frequently, their injuries are due to unexpected bugs in the software. Less commonly, injuries can occur during the learning procedure. However, they can teach and test robots at lower operating speeds and hence reduce the likelihood of injury.

C. Classification of injuries

In [3] injuries are classified according to their type between pinch (56%) and impact (44%) injuries. Pinch injury occurs when a robot traps a worker between itself and an object, whereas impact injury occurs when robot and worker collide. Consequences are classified as minor with no lost work-time, lost work-time injuries and fatal injuries. Based on results shown in [3], pinch injuries seem to be of a more serious nature than impact accidents.

A more recent approach gives a more detailed classification of contact scenarios that could lead to an accident [4]. Contact can be an unconstrained impact, partially constrained impact, constrained impact, clamping in the robot structure or secondary impact. Fig. 2 depicts these scenarios. Each of these impacts can be with a blunt or a sharp surface.

The International Organization for Standardization gives in [5] an extensive classification of hazards based on their origin. Mechanical hazards arise from unexpected or unintended

movements, unintended release of tools, rotational motion, trapping of clothes or hair, becoming trapped within the robot cell, etc. Electrical hazards comprise for instance contacts with live parts or connections, or exposure to arc flash. Thermal hazards are associated with hot surfaces or exposure to extreme temperatures required by an ongoing industrial process. Noise hazards are caused by loss of balance, disorientation or inability to coordinate tasks through communication. Other common hazards include vibration, radiation, hazards created by the use of dangerous materials, hazards coming from the dangerous environment in which robot is used (e.g., while trying to avoid a sharp edge, a worker comes in contact with a hot surface), and some combination of these factors.

Hazards associated with robots are well recognized, but the sources of hazards are frequently unique to a particular robot system. Not all of the identified hazards apply to all robots, nor is the level of risk associated with a given hazardous situation the same for each robot. Regulations defined by standards are thus ambiguous and not easily applicable. Designing measures that are robot and context specific may be one way of ensuring safety in constrained settings, such as industrial settings, but it falls short of ensuring safety in HRI in a generic way. This makes the problem of ensuring safety in human-robot interactions difficult.

III. INDUSTRIAL ROBOTS

Industrial robots were introduced in order to replace human workers performing dangerous, difficult, dull, monotonous and dirty tasks. In the past, these dangerous workplaces caused human workers injury and disease. Some health hazards that affect human workers in workspaces are toxic fumes, heat, radiation, noise, physical injuries and so on. In automated production systems, robots are deployed in large numbers for assembly, handling, welding and coating tasks. Robots thereby not only improve safety, but also productivity in heavy industry. However, as described in the previous section, robots can also represent a hazard for the people surrounding them. This section emphasizes accidents caused by robots in industry and gives an overview of safety measures proposed by current standards.

There are many standardization bodies dealing with the safety in human-robot interactions, but the most influential ones are the International Organization for Standardization (ISO) and the American National Standards Institute (ANSI). ISO 10218 is the most recent set of standards regarding the safety in industrial robots [5], [6]. The European Union has adopted it without any changes. The American standards manual ANSI/RIA R15.06-1999 was made public in 1999 [7]. It is widely used both in the United States and abroad (e.g., it has been translated to Japanese and used there). It was reaffirmed in 2009 without changes. Its newest version, ANSI/RIA R15.06-2012 is currently undergoing public review and will be made available soon [8].

Accidents related to industrial robots happen most frequently when the human worker is inside the work cell while the robot is operating. Additionally, incidents can occur during

servicing or programming the robot. Ideally, robots should be safe to humans regardless of failure or even misuse. In reality, perfect safety records in all contingencies are not feasible for machines that have to deliver performance in terms of cutting, welding, weight lifting, etc. A trade-off between performance and safety is therefore necessary.

One of the commonly used safety measures in a robot work cell is constructing a physical safety barrier around the robot. The main purpose of building the barrier is to safeguard the work cell against access by humans, while the robot is in operation. Safeguards can be fixed (e.g., safety fences) or movable (e.g., gates, flaps). Physical safeguards are complemented by non-physical safeguards (e.g., light curtains, scanners and safety mats). The American National Standard for safety requirements in industrial robots explicitly states that the most effective mean of safeguarding against injury is to shut the robot system off upon human entry in the robot work cell [9]. A safety fence usually has one or more gates attached to it, through which humans enter the robot work cell exclusively. An interlock device is incorporated into the gate, which stops the robot operation if the gate is opened. Closure of the gate does not resume the robot operation automatically. In accordance with [6], automatic operation shall be initiated from outside the safeguarded space and it should only be possible when all associated safeguards are active.

Means shall be provided to establish a restricted space around the robot. When it is not feasible to guard the robot work cell, limiting measures shall be used. The robot shall either have mechanical and electro-mechanical axis limiting devices, or software-defined limits to robot motion. Software-defined limits define a geometric shape (i.e., zone) that is used to either limit the motion of the robot within its interior, or limit the motion from entering this zone [5].

Danger to the human workers can further be reduced by proper installation of a robot system. By building high floor surfaces for covering the cables, the likelihood of tripping and falling over them is minimized. Restricted and operating spaces shall be established and clearly marked, as well as traffic routes (e.g., pedestrian aisles, visitor routes, etc.). Access and safe pathway to support services (electricity, gas, water), control systems, service and cleaning shall be provided [6].

Special attention needs to be devoted to the recovery from a failure. Loss of power or variations in power shall not result in a hazard. Re-initiation of power must not lead to any motion, as defined in [5]. Upon recovery, robot operation shall be re-initiated manually. Start and restart of the robot system shall be simple operations, and shall require relevant safety and protective measures to be functional. Location of actuating controls shall be chosen carefully, so as to prevent unintended operation. Status of actuating controls shall be clearly indicated (e.g., power on, fault detected, etc.).

In many cases, a collaboration between worker and robot is required. Therefore, it is not always possible to shut the robot off. A good solution is to equip the robot with force-torque sensor along with a force-torque control techniques, as described in [10], [11]. The force-torque sensor provides the amount and direction of the force and the moment, which

is then used to control the robot's movement considering its maximum allowable velocity.

In addition to safety control through force-torque sensors, the robot shall operate at slow speed mode when a human is present in a robot cell. ISO 10218 states that safe slow speed for a robot needs to be limited to $0.25 \text{ m}\cdot\text{s}^{-1}$. However, the size, configuration and environment of the robot to be guarded change the effectiveness of the safety techniques. Under certain environmental conditions or applications (such as programming or teaching), safe slow speed can be as low as $0.1 \text{ m}\cdot\text{s}^{-1}$ to $0.2 \text{ m}\cdot\text{s}^{-1}$ [12]. On the other hand, the authors of [4] criticize the limit of $0.25 \text{ m}\cdot\text{s}^{-1}$, saying that it is unnecessarily restrictive and as such greatly limits the robots' performance. In these two examples, the speed limit proposed by the standard does not generalize well to the variety of applications and scenarios. Hence, standards might be too loose and too general to apply easily, or they might require technical expertise beyond the user's capability [13]. To further demonstrate this difficulty, and also illustrate the previously mentioned software-defined limits, we refer to Section 5.12.3 from ISO 10218-1 [5]. It states that once set, safety-rated software limits shall always become activated upon robot power-up. There is, however, a possibility that a person installing a robot does not abide by the installation manual, and thus the software limits never get set.

Although the safe speed limits exist, human proximity to the robot can still be dangerous. Motions that pass near singularities can produce high axis speeds, despite the limits. These high speeds can be unexpected by an operator. Singularities should be avoided if possible. If not, the robot motion should be stopped and a warning should be provided prior to the robot passing through or correcting for a singularity. However, in some cases the singularities can be controlled without creating any hazardous motion [5].

Many approaches that endow the robot work cell with multiple sensors (e.g., proximity sensors, area detectors, cameras) to help determine that a human entered the vicinity of the robot have been studied in the literature. For example, Kulic and Kroft studied approaches that use this information to replan on the fly a path that minimizes chances of collision [14]. There are also industrial products based on the same principle, such as ABB's SafeMove [15]. Despite these achievements, adjusting the plan of the robot on the fly remains a challenge. While there exist control methods to replan on-line trajectories, software of industrial robots is not flexible and relies on predefined trajectories. To endow industrial robots with flexibility will hence require a complete redesign of the software provided currently by the main robotics manufacturers.

Apart from ISO 10218, the new ANSI/RIA R.15.06-2012 regulations will also give specific guidelines for collaborative robot applications [8]. It will be the first time that the collaboration is allowed by the American standard since the ANSI/RIA R.15.06-1999 prohibited a man in the loop with a robot in 1999. As the technology advances in terms of safety sensors and robot control, it has allowed man to be reintroduced in the loop. It will however be an optional



Fig. 3. Human-robot-cooperation for handling tasks. Inside a regular robot cell which is secured by light curtains, the robot handles gear boxes at regular speed in fully automated mode. Upon a human approaching the light curtain, the robot goes in the reduced-speed mode. The worker grasps the safety switch which activates the force-torque sensor. The worker guides the robot almost effortlessly by its handle [16]. (Fraunhofer IPA, Stuttgart)

feature, only available on new robots and robot systems.

Fig. 3 shows one example of an industrial robot that complies to ISO 10218 international standard and implements some of the safety measures we described above. The robot shown there serves as a human precision augmentation tool, and stretches from fully automated operation to acting as a servo-controlled balancer. The example shows the heavy gear box, that is grasped by the robot and softly balanced, so that worker can insert it precisely in the housing [16].

There are more potential ways to protect the human operators in the work cell, while loading and unloading the work parts in the processing operations. During this process, a rotary table can be used to transport raw materials from the human workers to the robot and finished work parts back to the workers. A rotary table is advantageous in that it divides the working space of the worker and the robot, reducing the risk of collision with the robot.

Despite all the safety measures, accidents can still happen, when least expected. ISO 10218-2 defines that every robot shall have a protective stop function and an independent emergency stop function [6]. It should be located at the minimum (albeit safe) distance, such that the operator has a quick, unobstructed access to it. Upon collision, robot should perform a safety stop function and remain still until manually reactivated. In this way, the damage it causes will be minimized. However, it is not good practice to have the robot immediately start moving in the other direction, away from the person with whom it came in undesired contact. In that case, the robot could possibly collide with another person who might be nearby and was not expecting a sudden change in the robot's moving trajectory.

Determining adequate reactive behavior when the robot enters in contact with humans is one of the many challenges robotics will have to solve to ensure safety in dynamic environments inhabited with humans. A two-dimensional tactile sensor has the ability to detect pressures and their two-dimensional distribution. Contact sensing provides an important and direct feedback for control, both in cases of voluntary and involuntary interactions with the environment. It can be applied to a robot as an artificial skin and used to safely stop a robot's movements in a human-centric perspective [17].

Tactile sensor placement is very important. It is worth noting that approaches that put tactile sensors along the main segment of the robot do not bring more information than would be provided by a force/torque sensor placed at the joint. Having tactile sensing on parts of the robot segments that bend or retract would be more useful, as it may allow to detect if a human limb got hooked in-between two retractable parts of the robot.

IV. MOBILE ROBOTS

A variety of mobile robotics are rapidly emerging. Unlike the industrial robots, in this case there is neither a dedicated robot workspace, nor trained personnel that might need to enter that workspace in order to interact with the robot in one way or another. Physical contact between humans and robots is unavoidable and even desirable when it comes to service or educational robots. Mobile robots are being integrated into the human environment. It is a highly dynamic environment, where humans move unpredictably and perform their everyday activities. Mobile robots are meant to assist humans, without jeopardizing their safety and exposing them to any hazard risk. In general, this holds true, but robotic cars show a particular danger due to their high mass and high energy movement.

There are still no standards dedicated to defining the safety requirements in non-industrial robot settings. ISO 10218 could partially be applied, as it defines collaborative requirements between humans and robots. The International Organization for Standardization has also devoted a few working groups to developing standards for other settings, such as robots in personal care, mobile service robots and medical robots. The new ISO/TS 15066 (scheduled to be released in 2013) should specify force and pressure limits for a safe contact between human and robot [8]. One of the goals of the new standard is to define metrics for expressing the relationship between safety and productivity.

A. Mobile Robots in Dynamic Environments

In dynamic environments, safety needs to be guaranteed without restricting robot system autonomy and flexibility. Therefore, not many principles from industrial robots' safety standards are applicable. Robots deployed in dynamic environments rely heavily on their sensory system, which provides them with perception of the environment around them. Robots physically embody the link between perception and action.

In many applications, robot precision is of crucial importance, especially when discussing safety. Because no perfect sensor exists, sensor readings are always accompanied by measurement noise. Sometimes, sensor data can be misleading, and if a robot entirely relies on it, it can consequently create a hazardous situation. Although information fusion is a very old problem, it has become a very popular subject in the last years. Multiple sensor data fusion is a process which combines observations from multiple sources, with the goal of providing robust and complete description of the environment [18]. Sensor data is usually combined by means of probabilistic filtering, which takes into account uncertainties associated with each sensor (cf. [19]).

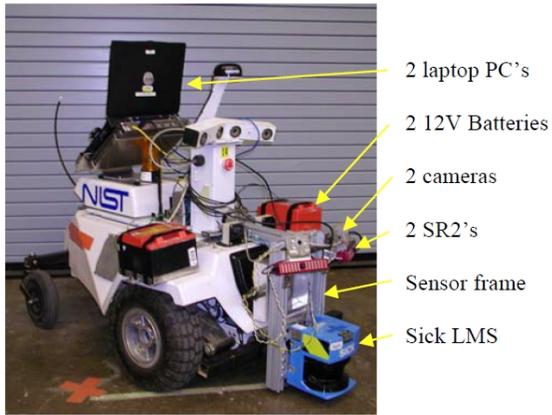


Fig. 4. DARPA LAGR vehicle equipped with stock dual stereo cameras, two infrared sensors, laser range-finder, and a physical bumper along with two drive-wheel encoders, an inertial measurement unit (IMU) and a global positioning system (GPS) sensor [21].

Mobile robots can be equipped with different kinds of sensors, such as ranging sensors (e.g., sonar, radar, laser, capacitive), vision-based sensors, tactile sensors, wheel and motor encoders, orientation and heading sensors, beacon-based position sensors, etc. Two or more sensors can measure the same physical quantity (e.g., distance from some object), or they can be complementary (i.e., measuring different quantities, or pointing in different directions, etc.). In either case, data coming from sensors can be fused, in order to get an integral perception of the environment and increase confidence in sensed data (for definition of confidence, cf. [20]). An example of a mobile robot equipped with multiple sensors is given in Fig. 4.

Information about the environment is then used for localization and planning purposes. One of the safety concepts is based on eliminating contact between robots and other objects (and especially humans). Using this as a limiting parameter, sensor outputs are used to plan trajectories that avoid obstacles, and detect other moving objects and avoid collision with them as well. Optical sensors can track people and their movements, compute dynamically safe zones and adjust a robot's speed and direction of movement to the given situation. One possible solution was proposed in [22]. It consists of two conventional cameras equipped with fisheye lenses, which can cover the complete robot workspace and provide a computationally effective alternative to a larger set of cameras or rotational cameras. The two cameras work in combination with a sophisticated background maintenance technique, which enables the detection, identification and tracking of moving objects and of humans with which the robot might need to interact. This solution is also robust to sudden changes in illumination.

In multiple robot settings, communication between robots can further enhance safety. A robot can exchange locally available information with other robots in its environment, which might help it to make an intelligent and safe decision about its action. For example, due to an obstructed view by another robot, the first robot might overlook an approaching dynamic object and in turn make an unsafe decision. If the

two robots exchange their sensor data among themselves, it is possible to gain a more complete and accurate situational awareness through the results of a sensor data fusion process [23]. Furthermore, a swarm of robots that mutually communicate can form a system whose goal is to improve safety in human-robot interaction. As shown in [24], a swarm forming an intelligent sensor system could be placed inside a piece of machinery to work as an inspection tool, increasing the safety of machinery with respect to humans, as well as limiting its downtime.

Care robots represent another growing market segment. Whereas ISO 10218 targets the robots for industrial environments and collaborative operations of human workers with robots, requirements for personal care robots have not been considered. A new working group is currently approved to formulate a new international safety standard for non-medical personal care robots [25]. The new standard under the name of ISO/DIS 13482 will deal with mobile servant robots, people carrier robots, physical assistant robots, personal care robot devices, and a few more. It shall provide risk assessment, as well as risk reduction or elimination and testing methodology for safety engineering procedures. At the time this survey concluded, it was still under development.

B. Robotic Cars

Robotics and intelligent systems are getting more and more involved in the automotive industry. Probably the most important goal for their involvement is improving the safety of passengers and other people, as well as increasing the performance of the vehicles. However, they are also introducing new types of hazards, and sometimes present a hazard themselves. In this section we give an overview of how robots improve and decrease the safety in transportation.

We have been witnesses of many car accidents and fatalities arising from them. In 2010, 32,778 people died in motor vehicle traffic crashes in the United States, according to the U.S. Department of Transportation [26]. Statistical projection showed an estimated number of 32,310 fatalities in 2011 [27], which would represent a decline of about 1.7 percent as compared to 2010. That is the lowest number of fatalities registered since 1949 (30,246 fatalities in 1949 [26]).

In order to reduce the number of accidents, there has been a lot of research conducted in the field of driver assistance. Among the most investigated fields are the braking assistance and collision avoidance systems. This is easily explicable based upon human behavior. Specifically, humans easily underestimate the speed of their vehicle and the distance to the vehicle in front of them. Since the vehicles operate at high speeds, and have enormous kinetic energy, injuries arising from car accidents are often fatal. For instance, an inattentive driver might overlook a stopped or a slowly moving vehicle ahead, or underestimate its threat level until it is too late. Another likely scenario is one of the major causes of accidents on the highways. Essentially, in this scenario, the distance to the vehicle ahead of the driver is not very large (e.g., less than 50 m) at the moment when the front vehicle suddenly starts to brake. The requirement for the driver of the

rear vehicle to detect an abrupt change in relative distance and acceleration is high and this often leads to accidents due to lack of attention, tiredness, fatigue, etc. These are just some of situations where the collision avoidance and braking assistance systems could be of crucial importance.

A number of measures for threat assessment in collision avoidance systems have been proposed throughout the literature. One common measure is *time-to-collision*, which represents the time it would take for a collision to happen at the current speed, distance and trajectory associated with the closest vehicle ahead [28], [29]. A second common time-based measure is a *time headway*. It specifies how much time the driver has to brake in case the car ahead suddenly brakes at its maximum deceleration level [30]. A more recent measure, *time-to-last-second-braking* [31], quantifies the danger or threat level more objectively and assesses the urgency level for the required protective action, e.g. braking. Related distance-based measures also exist, such as *projected minimum distance* during a collision avoidance process [32]. An important deceleration-based measure is *required deceleration*, which measures a constant deceleration level required to avoid a read-end collision [33].

Many warning and overriding algorithms have been developed based on the threat assessment measures. Their goal is to issue a threat warning to the driver at the right time, or to issue an overriding command (e.g., automatic braking) if there is no more time left for the human driver to react, following human reaction times [34]. Lee and Peng explain in [35] the algorithms proposed from researchers at Mazda, Honda, JHU and Jaguar, and give their subsequent evaluations. The evaluation results show that these algorithms are still far from being perfect—due to their having a low true positive rate, they are often seen by drivers as a nuisance. For these reasons, authors of [31] claim that warning and overriding criteria based on *time-to-last-second-braking* measure provide more appropriate warning and more effective overriding at the last moment.

Recent research takes a different approach to the same problem. Cabrera et al. propose in [36] a *time-to-last-second-acceleration* algorithm, that represents a warning algorithm for the driver of the front vehicle. It computes the time left before the extreme evasive action (i.e., full throttle acceleration) needs to be performed by the driver of the front vehicle, in order to avoid the rear-end collision.

These algorithms are directly related to the safety of the human driver and the passengers in the vehicle. For vehicles that have the above-mentioned collision avoidance systems enabled, we may say that they are semi-autonomous, since they have an automatically controlled braking and stopping system. Braking decisions are intelligently made and are based on the car's sensor readings. Hence, this is an exact application of robotic systems for improving human safety.

Whereas most of the automotive vendors focus on collision avoidance and adaptive cruise control systems (which basically follow a car ahead at a given distance), most of the academic laboratories that are in the field of intelligent vehicles, as well as Google, seek to produce a totally

autonomous system. There are many people in the field who claim that autonomous vehicles are going to decrease the number of accidents and eventually perform better than human drivers [37], [38]. For instance, whereas humans have the ability to track only a few objects in the environment at once, robots can do much more. Having a software that can track one object, tracking many objects simultaneously is not much harder. In fact, the autonomous car and the software on the car are, in the long run, potentially better than people in keeping track of a lot of elements going on in the environment [37].

However, having cars that drive themselves autonomously represents a new level of threat that previously did not exist, not just for the passengers in the car, but also for the whole environment—and most importantly, humans in the surrounding environment. Cars operate at high speeds, so the accidents in which they are involved often end up with casualties.

State-of-the-art autonomous vehicles are equipped with a number of sensors, including radars, laser range-finders, cameras, proximity sensors, global positioning systems (GPS), inertial navigation systems [39], etc. They use sensory information to localize themselves in the environment and make distinctions between drivable surfaces and obstacles. Sensors need to enable vehicles to *see* far ahead (approx. 150-200 m), so as to be able to drive fast and adjust their speed in accordance with the situation on the road in good time. Environments in which autonomous vehicles operate are highly dynamic. Many objects are not found in maps and cannot be modeled with respect to sensors output, as their presence is highly unpredictable.

There are situations when not even humans can avoid accidents (cf. [26] for statistics). Sometimes there is just not enough time to brake. Take as an example a child chasing a ball into the street. Even when assuming the robotic car was obeying all traffic rules and was actively aware of the surroundings, upon the kid or the ball appearing from behind a fence, bush, car or other obstruction, an accident could simply be unavoidable. However, since human reaction time is quite long, under identical circumstances, the child might have better chances to stay safe and get its ball if behind the wheel was not a human driver, but a robot. In fact, current robotic cars have more sensors and faster reaction times that could potentially vastly outperform a human driver.

Bicycles also present a safety issue. In the worst case situation, a cyclist can fall off their bike and stop immediately, lying in the road. A vehicle following a bicycle has to leave enough space to assure it can stop before such an occurrence, including reaction time. As already mentioned, reaction time should be better for robotic cars than for humans. Humans generally do not leave enough space. They leave even less space behind cars because cars actually cannot stop instantaneously, and they brake as the other cars brake. If a car hits another car at slow speeds, it is *tolerable*—nobody will be seriously hurt. Hitting a cyclist or pedestrian at slow speeds can mean death.

It is a real challenge to find a good balance between colliding into someone a car might hurt and abrupt braking because of an *obstacle* as small as a butterfly. No matter how

low a rate of accidents robotic cars manage to achieve, there will always be their opponents. Questions of liability will arise—is it the fault of the car owner, vehicle manufacturer or even programmer who worked on a software? As a matter of fact, the cause of accident might be easier to prove because of all the cameras and measurement equipment, which are able to record what is happening around the incident. A good long-term solution might be to have a black box installed in each autonomous vehicle, similar to airplanes. The effectiveness and liability of safety measures for autonomous vehicles will play a key role in the acceptance by society of robotic cars in the years to come. However, this and other ethical issues are beyond the scope of this survey.

The first license for driving an autonomous vehicle on American public roads was recently issued [40]. It symbolizes the safety plans set by Google, their employee training, system functions and accident reporting mechanisms. It is, despite the autonomy of the vehicle, a testing license, which demands the presence of persons in the vehicle while testing it on the public roads. This approval represents the first stepping stone towards a massive involvement of autonomous vehicles in traffic. These robots will certainly represent a phenomenon that will get a lot of attention within the robotic safety community in the years to come.

C. Miniature Mobile Robots

Based on a totally different principle, miniature robots have certainly contributed to safety in human-robot interaction. Haddadin et al. emphasize in [41] the impact of a robot's mass and velocity on the injury level. Their conclusion is that the best way to reduce the injury potential is by reducing a robot's weight. Mondada et al. recognized the potential of miniature mobile robots almost two decades ago [42]. Miniature robots had a great impact on control and evolutionary algorithm investigation over the last two decades. Bigger and heavier robots are not very safe for the investigation of unreliable algorithms. Because of their mass and moving force, the damage they can cause to the objects around them, especially human beings, can be considerable. Experiments with miniature robots can be performed in a small arena safely and cost-effectively. If the control over a miniature robot is lost, the robot will just end up in one of the arena walls, and refrain from damaging anything in the environment or itself.

The fact that most service and entertainment robots present in home environment are small and not powerful has ensured that these interactions are practically safe. However, miniaturization of the robot is not the ultimate solution. For many jobs dedicated to robots, a robot's weight and size is crucial. Many robots need to apply certain (high) force in order to perform their operations. In these scenarios, miniature robots are not applicable.

D. Assistive Devices

Approaches using robotic assistance for human rehabilitation, motor learning and recovery are common in the literature. Robots can assist, enhance, evaluate and document orthopedic

and neurological rehabilitation of movements. Rehabilitation devices for supporting rehabilitation of arm motion after stroke [43] and of locomotion after spinal surgery [44] are among the most advanced platforms. Commercial solutions, such as Lokomat, are successful examples of commercial assistive devices used in regular therapy sessions. While the forces that these robots can apply remain small, hence limiting greatly the risk of damaging the assisted limb, safety issues are a primary concern when designing wearable devices at large (including the design of devices to support or extend human senses). Limiting the forces does not guarantee safety as it is difficult to determine what is the minimal amount of force/torque that the device can safely apply when it is so close to the human body, as sometimes a small force applied at a very weak point can lead to serious injury or death.

When designing assistive devices, safety issues arise from the use of shared control. If an accident happens, who of the machine or the human should be held responsible, since they both share the control? The interested reader can refer to [45] for a discussion on these issues. Consider wheelchairs driven by brain-machine interface [46]; unless one records all data input and output from the system (as proposed in Section IV-B), it will be difficult to determine whether the person gave the wrong command or the system misinterpreted the person's intention.

V. CONCLUSION

Robots are slowly and increasingly pervading in many segments of human lives. They are becoming part of our living environment. While useful, robots also represent a potential hazard. They can move their arms or bodies forcefully and very rapidly, and often manipulate dangerous and sharp tools. This represents a threat to all living agents that are surrounding robots. If humans are present in the robots' proximity, the situation gets even more dangerous.

Industrial robots are usually installed inside a guarded work cell and their operation is regulated with standards. Relaxation of some of the rules defined by standards makes the issue of safety a prominent one. Therefore, defining robot and context specific safety guidelines needs to be urgently addressed by the scientific and industrial community.

Mobile robots rely on multiple sensors and sophisticated control strategies in order to ensure safety in their interactions with humans. Although robot safety has long been an issue of concern, accidents arising from this type of interaction threaten to grow to a significant number with the growing number of mobile robots in the human environment.

Without any doubt, autonomous vehicles represent one of the areas that will have the biggest impact in the near future. Many experts in the field claim that autonomous vehicles are going to increase traffic safety. Still, they also represent a new threat. Therefore, there is an immediate need for the design of clear guidelines and of safety measures to ensure that these new set of autonomous and semi-autonomous vehicles increase safety on our roads, as opposed to the converse.

As robots become more complex, with more and more degrees of freedom (e.g., humanoid robots), the risk of trap-

ping in-between two joints will increase; higher complexity in robot's range of motion will make it more difficult for humans to predict the robot's displacement. To address the numerous challenges enumerated before requires i) the design of new sensing technology and of fast sensor fusion algorithms to track multiple moving targets in real time, ii) to achieve robust detection of human motion in order to build good predictive systems, iii) to ensure robust detection of contact between robots and surrounding living agents in multiple points, and iv) to develop fast responsive controllers that can replan trajectories in complex, cluttered environment in real time.

REFERENCES

- [1] O. Ogorodnikova, "Methodology of safety for a human robot interaction designing stage," pp. 452–457, May 2008.
- [2] J. Carlson and R. Murphy, "How UGVs physically fail in the field," *IEEE Transactions on Robotics*, vol. 21, pp. 423–437, June 2005.
- [3] B. C. Jiang and C. A. Gainer Jr., "A cause-and-effect analysis of robot accidents," *Journal of Occupational Accidents*, vol. 9, pp. 27–45, June 1987.
- [4] S. Haddadin, A. Albu-Schffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *The International Journal of Robotics Research*, vol. 28, pp. 1507–1527, Nov. 2009.
- [5] ISO, *ISO 10218-1:2011: Robots and robotic devices Safety requirements for industrial robots Part 1: Robots*. Geneva, Switzerland: International Organization for Standardization, 2011.
- [6] ISO, *ISO 10218-2:2011: Robots and robotic devices Safety requirements for industrial robots Part 2: Robot Systems and Integration*. Geneva, Switzerland: International Organization for Standardization, 2011.
- [7] ANSI/RIA, *ANSI R15.06-1999: American National Standard for Industrial Robots and Robot Systems—Safety Requirements*. American National Standards Institute/Robotics Industry Association, 1999.
- [8] J. Fryman, "Robotic industries association webinar: 2012 ANSI/RIA robot safety standard preview," Aug. 2012.
- [9] ANSI/RIA, *ANSI R15.06-1992: American National Standard for Industrial Robots and Robot Systems—Safety Requirements*. American National Standards Institute/Robotics Industry Association, 1992.
- [10] J. N. Pires, J. Ramming, S. Rauch, and R. Arajo, "Force/torque sensing applied to industrial robotic deburring," *Sensor Review*, vol. 22, no. 3, pp. 232–241, 2002.
- [11] S. Kuhn, T. Gecks, and D. Henrich, "Velocity control for safe robot guidance based on fused vision and force/torque data," pp. 485–492, Sept. 2006.
- [12] W. Karwowski, H. Parsaei, A. Soundararajan, and N. Pongpatanasuegsa, "Estimation of safe distance from the robot arm as a guide for limiting slow speed of robot motions," *HUMAN FACTORS SOCIETY, P. O. BOX 1369, SANTA MONICA, CA 90406-1369(USA)*, 1992.
- [13] R. A. Hirschfeld, F. Aghazadeh, and R. C. Chapleski, "Survey of robot safety in industry," *International Journal of Human Factors in Manufacturing*, vol. 3, pp. 369–379, Mar. 2007.
- [14] D. Kulic and E. A. Croft, "Safe planning for human-robot interaction," *Journal of Robotic Systems*, vol. 22, no. 7, p. 383396, 2005.
- [15] C. Thomas, F. Busch, B. Kuhlencoetter, and J. Deuse, "Ensuring human safety with offline simulation and real-time workspace surveillance to develop a hybrid robot assistance system for welding of assemblies," in *Enabling Manufacturing Competitiveness and Economic Sustainability* (H. A. ElMaraghy, ed.), pp. 464–470, Springer Berlin Heidelberg, Jan. 2012.
- [16] M. Hgele, K. Nilsson, and J. N. Pires, "Industrial robotics," in *Springer Handbook of Robotics* (B. Siciliano and O. Khatib, eds.), pp. 963–986, Springer Berlin Heidelberg, 2008.
- [17] B. D. Argall and A. G. Billard, "A survey of tactile HumanRobot interactions," *Robotics and Autonomous Systems*, vol. 58, pp. 1159–1176, Oct. 2010.
- [18] H. Durrant-Whyte and T. C. Henderson, "Multisensor data fusion," in *Springer Handbook of Robotics* (B. Siciliano and O. Khatib, eds.), pp. 585–610, Springer Berlin Heidelberg, 2008.
- [19] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents, Mit Press, 2005.
- [20] M. Siegel and H. Wu, "Confidence fusion [sensor fusion]," pp. 94–99, May 2004.
- [21] R. Bostelman, T. Hong, and T. Y. Chang, "Safety standard advancement toward mobile robot use near humans," (Chicago, IL, US), Sept. 2005.
- [22] E. Martinez and A. del Pobil, "Safety for human-robot interaction in dynamic environments," pp. 327–332, Nov. 2009.
- [23] Car to car communication consortium, see <http://www.car-to-car.org>.
- [24] N. Correll and A. Martinoli, "Multirobot inspection of industrial machinery," *IEEE Robotics & Automation Magazine*, vol. 16, pp. 103–112, Mar. 2009.
- [25] Y. Yamada and Y. Ota, "Novel activity on international safety standardization for personal care robots," pp. 1882–1883, Aug. 2009.
- [26] U. D. of Transportation, "2010 motor vehicle crashes: Overview," Research Note DOT HS 811 552, National Highway Traffic Safety Administration, Feb. 2012.
- [27] U. D. of Transportation, "Early estimate of motor vehicle traffic fatalities in 2011," Research Note DOT HS 811 604, National Highway Traffic Safety Administration, May 2012.
- [28] A. Van Der Horst, "Time-based analysis of road user behavior in normal and critical encounters," Apr. 1990.
- [29] R. J. Kiefer, M. T. Cassar, C. A. Flannagan, D. J. LeBlanc, M. D. Palmer, R. K. Deering, and M. A. Shulman, "Refining the CAMP crash alert timing approach by examining last-second braking and lane change maneuvers under various kinematic conditions," Final Report DOT HS 809 574, CAMP, NHTSA, Jan. 2003.
- [30] R. G. C. Fuller, "Determinants of time headway adopted by truck drivers," *Ergonomics*, vol. 24, no. 6, pp. 463–474, 1981.
- [31] Y. Zhang, E. Antonsson, and K. Grote, "A new threat assessment measure for collision avoidance systems," pp. 968–975, Sept. 2006.
- [32] R. J. Kiefer, D. J. LeBlanc, M. D. Palmer, J. Salinger, R. K. Deering, and M. A. Shulman, "Development and validation of functional definitions and evaluation procedures for collision warning/avoidance system," Final Report DOT HS 808 964, CAMP, NHTSA, Aug. 1999.
- [33] S. J. Brunson, E. M. Kyle, N. C. Phamdo, and G. R. Preziotti, "Alert algorithm development program NHTSA rear-end collision alert algorithm," Final Report DOT HS 809 526, The Johns Hopkins University, Applied Physics Laboratory, Sept. 2002.
- [34] M. Green, "How long does it take to stop?" methodological analysis of driver perception-brake times," *Transportation Human Factors*, vol. 2, no. 3, pp. 195–216, 2000.
- [35] K. Lee and H. Peng, "Evaluation of automotive forward collision warning and collision avoidance algorithms," *Vehicle system dynamics*, vol. 43, no. 10, pp. 735–751, 2005.
- [36] A. Cabrera, S. Goyal, and A. Martinoli, "A new collision warning system for lead vehicles in rear-end collisions," (Alcala de Henares, Spain), June 2012.
- [37] D. Huttenlocher, "Robots podcast: Cornell racing team and velodynes LIDAR sensor, episode 1, June 2008."
- [38] S. Thrun, "Udacity CS373: Artificial Intelligence, Programming a Robotic Car, Unit 7, Test Drive," Feb. 2012.
- [39] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhne, D. Johnston, S. Klumpp, D. Langer, A. Levandowski, J. Levinson, J. Marcil, D. Orenstein, J. Paefgen, I. Penny, A. Petrovskaya, M. Pflueger, G. Stanek, D. Stavens, A. Vogt, and S. Thrun, "Junior: The stanford entry in the urban challenge," *Journal of Field Robotics*, vol. 25, pp. 569–597, Sept. 2008.
- [40] N. D. of Motor Vehicles, "Nevada DMV issues first autonomous vehicle testing license to google," May 2012.
- [41] S. Haddadin, A. Albu-Schaffer, M. Frommberger, and G. Hirzinger, "The role of the robot mass and velocity in physical human-robot interaction - part II: constrained blunt impacts," pp. 1339–1345, May 2008.
- [42] F. Mondada, E. Franzi, and P. Jenne, "Mobile robot miniaturisation: A tool for investigation in control algorithms," in *Experimental Robotics III* (T. Yoshikawa and F. Miyazaki, eds.), vol. 200 of *Lecture Notes in Control and Information Sciences*, pp. 501–513, Springer Berlin / Heidelberg, 1994.
- [43] T. Nef and R. Riener, "ARMin - design of a novel arm rehabilitation robot," in *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005*, pp. 57–60, July 2005.
- [44] Y. Stauffer, Y. Allemand, M. Bouri, J. Fournier, R. Clavel, P. Metrailler, R. Brodard, and F. Reynard, "The WalkTrainer—A new generation of walking reeducation device combining orthoses and muscle stimulation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, pp. 38–45, Feb. 2009.
- [45] P.-A. Mudry, S. Degallier, and A. Billard, "On the influence of symbols and myths in the responsibility ascription problem in roboethics - a roboticist's perspective," in *The 17th IEEE International Symposium on Robot and Human Interactive Communication, 2008. RO-MAN 2008*, pp. 563–568, Aug. 2008.
- [46] T. E. Carlson and J. d. R. Milln, "The robotic architecture of an asynchronous BrainActuated wheelchair," *IEEE Robotics and Automation Magazine*, 2013.