

On making robots understand safety: Embedding injury knowledge into control

The International Journal of
Robotics Research
31(13) 1578–1815
© The Author(s) 2012
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0278364912462256
ijr.sagepub.com



Sami Haddadin¹, Simon Haddadin^{1,2}, Augusto Khoury^{1,3}, Tim Rokahr¹, Sven Parusel¹,
Rainer Burgkart², Antonio Bicchi³ and Alin Albu-Schäffer¹

Abstract

Enabling robots to safely interact with humans is an essential goal of robotics research. The developments achieved over recent years in mechanical design and control made it possible to have active cooperation between humans and robots in rather complex situations. For this, safe robot behavior even under worst-case situations is crucial and forms also a basis for higher-level decisional aspects. For quantifying what safe behavior really means, the definition of injury, as well as understanding its general dynamics, are essential. This insight can then be applied to design and control robots such that injury due to robot–human impacts is explicitly taken into account. In this paper we approach the problem from a medical injury analysis point of view in order to formulate the relation between robot mass, velocity, impact geometry and resulting injury qualified in medical terms. We transform these insights into processable representations and propose a motion supervisor that utilizes injury knowledge for generating safe robot motions. The algorithm takes into account the reflected inertia, velocity, and geometry at possible impact locations. The proposed framework forms a basis for generating truly safe velocity bounds that explicitly consider the dynamic properties of the manipulator and human injury.

Keywords

Physical human-robot interaction, safety, biomechanics, human injury, soft-tissue, nonlinear control, impact dynamics, robot standardization

1. Introduction

Close physical human–robot interaction (pHRI) is one of the grand challenges of present-day robotics research. Several robotic systems gained sufficient control capabilities to perform delicate and complex manipulation and pHRI tasks that require the dynamic exchange of forces between the robot and its environment (Townsend and Salisbury, 1993; Albu-Schäffer et al., 2007a; Shin et al., 2008; Iwata and Sugano, 2009; Bischoff et al., 2010). This step made it possible, e.g., to automate difficult and up to now still manually executed assembly tasks. In particular, the achieved sensitive and fast manipulation capabilities (Hogan, 1985; Goldsmith et al., 1999; De Luca and Mattone, 2004; Zollo et al., 2005; Albu-Schäffer et al., 2007b; Stemmer et al., 2007; Haddadin et al., 2008) of these robots prevent damage from the potentially fragile objects handled and are less dangerous to humans that may directly interact with the device. In order to enable such direct physical cooperation between human and robot, there has been strong interest recently in removing classical safety barriers, such as fences or light curtains for novel ‘human-friendly’ robots made for direct interaction, see Figure 1.

1.1. State of the art in robot safety analysis

In order to define the regulations needed, one has to understand what it means to design and control robots such that they operate safely in human environments. The most stringent requirement is therefore to ensure that a human would not suffer any severe injury even under worst-case conditions. Based on the pioneering work of Yamada (Yamada et al., 1996, 1997), where human pain tolerance was introduced as a criterion for safe robot impact behavior, others have extended this line of research considerably in terms of novel actuation mechanisms, interaction control

¹Robotics and Mechatronics Center, DLR - German Aerospace Center, Wessling, Germany

²Orthopedic Clinic and Polyclinic, University Hospital *Klinikum Rechts der Isar*, Technical University of Munich (TUM), Munich, Germany

³Interdepartmental Research Center “E. Piaggio”, Faculty of Engineering, University of Pisa, Pisa, Italy

Corresponding author:

Sami Haddadin, Robotics and Mechatronics Center, DLR - German Aerospace Center, Muenchner Strasse 20, Wessling 82234, Germany.
Email: sami.haddadin@dlr.de



Fig. 1. Human and robot closely working together is one of the main goals of current robotics research. In particular, robotic co-workers in industrial environments have already been introduced to some real-world settings.

schemes, and robot–human collision injury models. Ikuta et al. (2003), Bicchi and Tonietti (2004), Zinn et al. (2005), Heinzmann and Zelinsky (2003), and Lim and Tanie (2000) proposed various control strategies and novel joint/system designs for making robots safer in case of a potential collision with the human. Ikuta et al. (2003) proposed also some definitions of risk indicators based on limit contact forces and inertial robot properties. The work of Kulic and Croft (2007) directly relates to that of Ikuta et al. (2003) in terms of defining robot design related safety indices. Furthermore, they associate safety to human-centric quantities such as head orientation or affective state. Bicchi and Tonietti (2004) and Zinn et al. (2005) adopted the head injury criterion (HIC), a biomechanical injury indicator for blunt head collisions that stems originally from the analysis of acceleration/deceleration tolerance of the human head (Versace, 1971), to the robotics community. Oberer and Schraft (2007) analyzed this and some other indicators via finite element modeling (FEM) crash-test dummy simulations for an industrial robot colliding with a side-crash-test dummy. Park and Song (2009) developed a blunt impact robot–human collision model for analyzing several severity indices such as HIC. Further analysis on the HIC and its role in robotics are discussed by Gao and Wampler (2009). An analysis of impact energy density as an indicator for contusion was evaluated by Povse et al. (2010) based on the work of Haddadin et al. (2007a). Wassink and Stramioli (2007) derived an impact model for Hertzian contact to investigate blunt impact stress and associated skin injury. Further work on danger indices can be found in the work of Ogrodnikova (2009). An attempt to deduce a structured description of pHRI scenarios is presented by Matthias et al. (2010). From the standardization side, the ISO 10218 was the first step towards developing international guidelines for collaborative robots sharing their workspace with humans (International Organization for Standardization, 2006). The work of BG/BGIA (2009), which will become a part of the

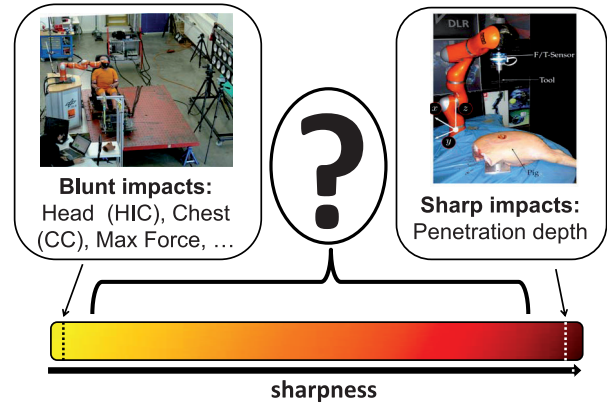


Fig. 2. Crash-testing and injury analysis in robotics was mainly concerned either with blunt impacts or very sharp contacts. However, the range of moderate sharpness is still a relatively open field.

technical specification of the new ISO 10218 version (International Organization for Standardization, 2011), occurring injuries in collaborative applications are required to cause maximally mild contusions.¹ This fact shows that the analysis and understanding of injury in robotics has become the essential prerequisite for real-world pHRI and the community should place great emphasis on it.

In the reviewed line of research, some of the earliest thorough safety studies over the last years were performed in Haddadin et al. (2007b, 2009b,a, 2010, 2011) and Park et al. (2011). This work led to insights into the potential injury a human would suffer due to a collision with a robot. There, various worst-case scenarios in pHRI were discussed and analyzed according to the following scheme:

1. select and/or define and classify the impact type;
2. select the appropriate injury measure(s);
3. evaluate the potential injury of the human;
4. quantify the influence of the relevant robot parameters;
5. evaluate the effectiveness of countermeasures for injury reduction and prevention.

This analysis revealed the basics of injuries that are caused, for example, by fast blunt impacts, dynamic and quasi-static clamping, or cuts and stabs by sharp tools.

1.2. Contributions

As described above, the research efforts on understanding safety were considerable over the last years, focusing mainly on blunt impacts, see Figure 2. However, as a robot is supposed to operate with different tools or grasp/use/carry sharp objects, the understanding of sharp soft-tissue injury is certainly an important aim. Sharp contacts and resulting soft-tissue injury, however, have up to now only been treated by Haddadin et al. (2010, 2011). This is mainly due to the lack of exploitable biomechanical data. The understanding of how impact mass, velocity, and contact geometry affect resulting injury, in general, has

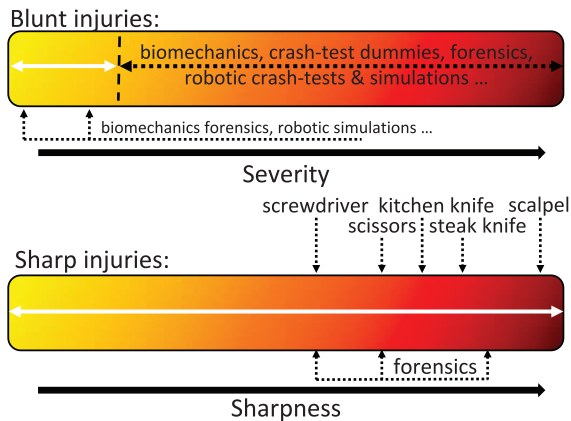


Fig. 3. Classification of existing work on safety in robotics and general biomechanics/forensics. Except for some punctual contributions, typical investigations regarding blunt injury were concerned with rather severe injury and left out low severity injury up to now. A full analysis of sharp injury is largely missing. Only very basic sample investigations were conducted so far.

not yet been systematically approached and only punctual contributions have been made in forensics and robotics, see Figure 3 (lower). Only Wassink and Stramigioli (2007) and Park et al. (2011) treated this topic to some extent for low-severity blunt injury.² The former focuses on a model for skin injury. The latter treats contusions and lacerations. These are predicted with a model for which the underlying data originates from existing cadaver experiments in the biomechanical literature, see Figure 3 (upper).

Another open problem in human-friendly robotics is how to embed the rather general understanding of injury in control. Usually, injury-related insights in robotics are either used to mechanically design safer robots, or to show that a particular design has beneficial impact characteristics. Taking injury knowledge explicitly into account was not done yet.

In this paper, we contribute to both aforementioned problems. First, we investigate the relation among impact mass, velocity, geometry, and medically observable soft-tissue injury by systematic drop-testing experiments with pig abdominal wall samples (Meyer, 1996). The occurring soft-tissue injury is assessed by means of the so-called AO classification (Rüedi et al., 2007), which is an international medical classification system that also involves concomitant soft-tissue injuries. Owing to the growing demand in standardization to focus especially on mild contusions, we also consider them in our analysis. These general insights are then used in an injury-based velocity controller for dynamically limiting the commanded velocity, taking into account the robot's instantaneous dynamic properties. In other words, we embed acquired 'injury knowledge' explicitly into the motion controller. Thus, no trajectories can be executed that would exceed certain maximum injury limits. Our approach to this problem is introduced hereafter in more detail.

1.3. Approach

The underlying motivation of this work is to make robotic systems safe for interaction without introducing across-the-board requirements, such as those originally demanded by ISO-10218-2006 or BG/BGIA (2009). All generated insights have very general character and are not tailored to a certain robot design. Our data, analysis, and control methodology can be used by the entire community and for every robot. As the focus of this paper lies on soft-tissue injury, we need to first understand how a generic collision between robot and human affects this. More specifically, one may ask how do impact mass, velocity, and contact geometry relate to injury. This information can then be used to design controllers that incorporate this knowledge explicitly, i.e. making robots aware of what safe interaction means on a very low-level already. This is especially important for letting them operate as fast as possible under the safety constraint, i.e. to give an answer to the question 'How fast can I move without hurting someone?'. Our previous work in particular gave answers to this question for blunt impacts and to some extent also for sharp soft-tissue injury potentially caused by the DLR Lightweight Robot III (LWR-III) that handles very dangerous tools such as knives and scalpels. However, before being able to understand general soft-tissue injury in robotics (in contrast to our previous work, where we mostly considered lethal injury due to sharp contact, we treat also mild injury in the present work), we first need sufficient biomechanical injury data for formulating the respective safety limits. As the systematic understanding of soft-tissue injury was not a major focus in biomechanics or in forensics, this makes it necessary to carry out impact tests for determining the relation between 'input' robot parameters (reflected inertia, velocity, and impact curvature) and the resulting injury. As we certainly cannot test every possible contact geometry, the first step was to identify relevant impact geometry primitives. Then, we carried out large-scale drop-testing experiments with fresh abdominal pig samples³ at varying mass and speed for a given primitive. The possibly produced injury is then medically assessed as follows:

1. immediate medical observation of the impact area;
2. dissection and injury analysis;
3. histopathologic evaluation.

The observed injury is then medically classified into standardized classes of injury in terms of the so-called AO classification. The overall approach is depicted in Figure 4. In this paper we analyze three selected primitives and one real-world object, requiring hundreds of drop tests.⁴ In order to make use of the functional relation ($robot\ parameters \rightarrow injury$) for robot control, we derive risk curves for every given primitive. This facilitates a simple and intuitive representation of the ($robot\ parameters \rightarrow injury$) relation, which can then be stored into a real-time database. The gathered knowledge is now accessible online for supervisory real-time velocity control. The proposed controller

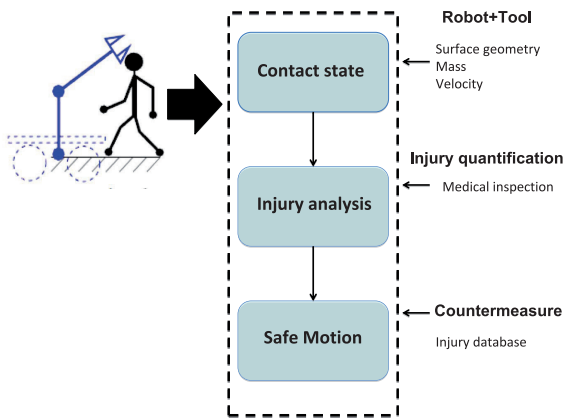


Fig. 4. Proposed methodology. Based on the mechanical input parameters mass, velocity, and surface geometry, a medical evaluation leads to the quantification of injury. The relation between them can then be used for a control strategy that ensures safe motion behavior.

takes into account the reflected dynamics of the robot at relevant operational points, their velocity, and respective contact geometry. This framework enforces safe impact characteristics on the robot in case of an unexpected collision with a human. It is important to note that our previous work on blunt impact analysis can be seamlessly integrated into the database-controller methodology.

One general remark regarding injury prediction shall be made at this point. For predicting injury one needs an appropriate model that reliably captures the occurrence of injury for a given mechanical input. Mostly, one would use the drop-test results (or any other collision experiment) for measuring impact **output** characteristics, such as contact forces and stress, and relate them to injury. In other words, one would, e.g., try to define ‘threshold forces’ or ‘threshold stress’ (possibly nonlinear functionals). In turn, these would be used to acquire contact models and predict the resulting injury via collision simulations, see Figure 5 (upper). Unfortunately, there are two major problems with such an approach. First, it is very difficult to measure quantities such as impact stress for complex geometries (which are, however, those we are interested in) and, second, the consistency with the medically observed injury is often insufficient. This can have multiple causes, of which the certainly most important one is that a single quantity cannot capture the complex behavior of human soft-tissue (especially in robot failure/injury cases). Therefore, we select a different approach, which uses medical observation of injury as ‘ground truth’, i.e. the independent variable, and derive risk curves relating physical **input** parameters (mass, velocity, and impact geometry) with injury.⁵ This approach is fundamentally different from previous approaches.

The remainder of the paper is organized as follows. Section 2 gives an overview of the soft-tissue injury mechanisms, penetration models of soft-tissue, and some mechanical properties of skin. In Section 3 drop-testing

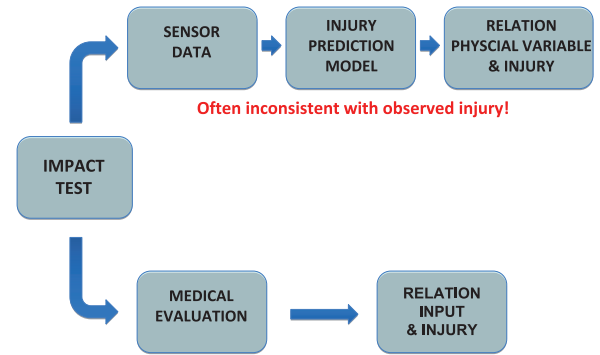


Fig. 5. Comparison of model-based approaches (upper path) versus medically oriented analysis (lower path) for injury analysis and prediction. Extracting models based on sensor data and mechanical human response for injury prediction is laborious and prone to false interpretation. In particular, the goal of relating certain physical quantities to injury seems to be very limited, as the human mechanical response cannot be fully described by such a simple model. The purely medically oriented approach substitutes the prediction with more general experimental investigations that cover the range of interest and directly relate medical findings with mechanical input parameters. This is the approach we take in the present paper.

experiments are presented that provide basic data for medical analysis in terms of the AO classification. Then, in Section 4 so-called ‘safety curves’ are used together with an accordingly designed injury database for velocity shaping based on the dynamic and surface properties of a manipulator. The controller is experimentally validated with the LWR-III. Section 5 concludes the paper and provides an outlook. Finally, the appendix contains several measurement results from the drop-testing experiments.

2. Basics of soft-tissue injury

Impacts can cause several characteristic soft-tissue injuries in the context of robotics. The resulting injury may be classified as follows (Edlich et al., 1988; Macchiarelli et al., 2002; Cainelli et al., 2008):

- bruises;
- abrasions;
- lacerations;
- puncture/stab wounds.

Their mechanisms are very different but some general characteristics can be found as described hereafter.

2.1. Mechanisms of soft-tissues injuries

Every traumatic wound is a result of a finite energy source that causes tissue disruption. Skin, fat, and muscles absorb the kinetic energy of an impactor during collision by compression and propagation of the stress waves into the body. The dynamics of this energy exchange determines the magnitude of lesion.

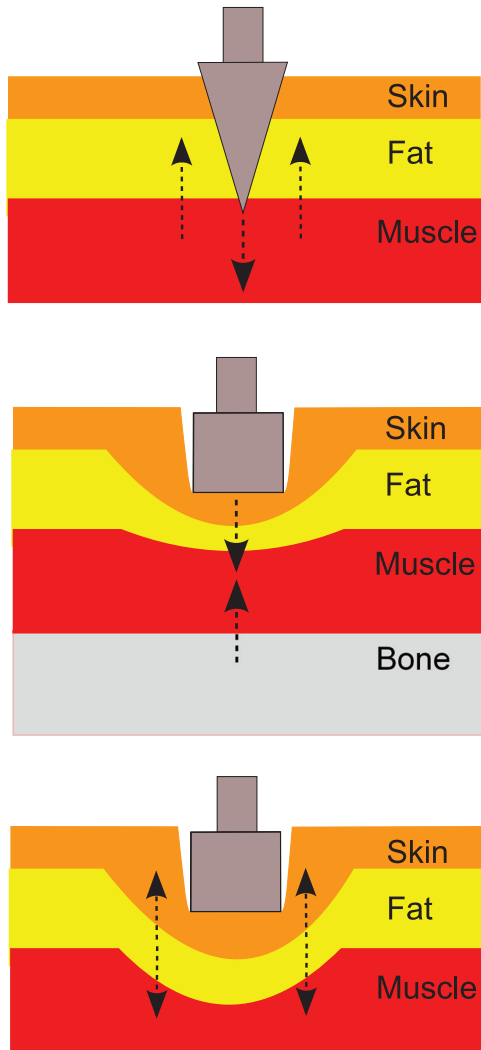


Fig. 6. Collision with a sharp impactor: injury is caused by shear forces (top). Collision with a blunt impactor: if the soft-tissue is supported by underlying bone skin, injury is due to compressive force (middle), otherwise the injury mechanism is tension (bottom), adapted from Shergold and Fleck (2004, 2005) to our cases.

Soft-tissue injuries can be caused by forces of different types. In the case of shearing-type hits with sharp objects, the applied forces counteract in parallel planes separated by a small distance, see Figure 6 (top). Shear forces are caused by the sharp edge of the impactor and by the reaction of the tissue itself. A linear laceration is a typical result of this trauma type. As the amount of tissue volume that is in contact with the sharp object is very small, only a very little total energy (< 100 J) is required to cause it.

If an injury is caused by a collision with a blunt object, the mechanism of injury can be compression or tension, see Figure 6 (middle and bottom). In both cases, forces act against the impactor, but unlike for the shearing case, they act within the same plane. If the blunt impactor hits soft-tissue that is not supported by the underlying bone

structure, the resulting injury mechanism is mainly tension. For soft-tissue supported by bone, skin injury is primarily due to compressive forces. The threshold energy for tissue failure in collisions with blunt impactors is considerably larger than that calculated for shear forces, as the energy is distributed among larger tissue volume. If the injury mechanism is compression, failure of bone-supported tissue occurs at energy levels of 2.52 J/cm^2 (Eisenmenger, 2004). This level of energy is comparable to that encountered by soft-tissue layers covering the cranium (weight ≈ 4 kg) during a car impact with a tree at 8 km/h, assuming the head hits the dashboard with an impact area of $\approx 8 \text{ cm}^2$.

2.2. Penetration models of soft-tissue

Unfortunately, existing literature provides only little insight into the underlying mechanisms of penetration. However, a limited number of experimental studies indicate that deep penetration of skin, rubber, and soft solid is characterized by a substantial reversible deformation and the start and growth of a crack. These studies also demonstrate that the tip geometry and material properties influence the shape of the crack in the penetrated solid. Results in Shergold and Fleck (2004, 2005) can be summarized as follows:

1. penetration force decreases with increasing sharpness of the penetrator;
2. penetration force is sensitive to the type of tissue being penetrated (for fat it is significantly lower than for skin and muscle);
3. penetration force depends on the degree of skin pre-stretch and penetrator velocity; and
4. tip geometry determines the penetration mechanism of a soft solid.

Following this general introduction into the mechanics of soft-tissue injury, we shortly review relevant pig anatomy hereafter, as we conduct our drop-testing experiments with pig subject. This review is important to understand our interpretations and consecutive implications.

2.3. Anatomy of the pig specimen

As in humans, pig skin consists of an epidermis, an underlying dermis, and subcutaneous adipose tissue (Ross and Pawlina, 2006), see Figure 7. In general, pig tissue can be regarded as a well suited substitute for human tissue (Meyer, 1996). In our experiments we used pig processed by the slaughterhouse, meaning that the epidermis was scalded off. Consequently, the specimen is lacking epidermis (making our analysis conservative). The dermal thickness was mostly < 2 mm, whereas the subcutaneous fat varied between 3 and 17 mm.

Beneath skin and fat lies a rather thick connective tissue sheath, separating it from the skeletal muscles. The human counterpart of this sheath would be fascia of Scarpa. The muscles of the anterolateral abdominal wall are a strong




Fig. 7. Layered structure of the specimen abdominal wall. The tissue consists of the skin, fasciae, muscles, fascia transversalis, and the peritoneum.




Fig. 8. The layered structure of the human skin. It can be roughly divided into epidermis, dermis, and subcutaneous tissue. In the latter, we can find venules and arterioles.

striated muscular wall, consisting of three flat muscles (external oblique abdominal m., internal oblique abdominal m., transverse abdominal m.), and two vertical muscles (rectus abdominis and pyramidal muscles). The flat muscles are crisscrossing each other and are covered by superficial, intermediate, and deep layers of investing fascia. The remaining parts of the abdominal wall are the muscular aponeurosis, which is the anterior and medial extension of the flat muscles' tendons, transverse (endoabdominal) fascia, extraperitoneal fat, and parietal peritoneum (Moore et al., 2009). Since these were rather inconsistently present in the considered specimen, we excluded their analysis from our experiments. If significant exceptions could be found (very obvious injuries to those structures) they were, however, considered.

Apart from the anatomy of pig tissue, the mechanical properties of mammalian skin are worthy of summarization especially for the sake of terminology.

2.4. Mechanical properties of skin

The mechanical properties of human skin such as uniaxial stress versus strain and toughness, are extensively

reviewed by Shergold and Fleck (2005) and Shergold et al. (2006). Here, we give only a very short summary. Skin consists of the two main tissue layers dermis and epidermis, see Figure 8. In mammalian skin the dermis is typically $20 \times$ thicker than the epidermis and characterizes the overall constitutive behavior. The major structural components of dermis are collagen fibers, accounting for 60–80% of its dry weight. Consequently, the constitutive behavior of skin depends upon the structure and density of collagen fibers found within the dermal layer. The greater compliance of skin, compared with an individual collagen fiber, is attributed to the capacity of the network fibers to straighten and align in the direction of applied strain.

In general, it is important to note that human tissue shows significant variability and several rather complicated aspects such as its hysteresis behavior (Fung, 1993). In general, the large variability in skin responses are hard to capture. Acquired data varies by individuals, subject age, impact location, and the method of storing the sample prior to testing. Furthermore, the orthotropic constitutive response of skin is an interesting behavior. In the dermal layer, the collagen fibers show preferential orientations. These directions are called Langer's lines and are the trajectory of maximum skin tension and least flexibility. Consequently, the onset of strain hardening begins at lower strains when the skin is stretched parallel to the Langer's lines, compared with stretching perpendicular to this direction.

After this introduction of soft-tissue injuries and the mechanics and anatomy of skin properties, we discuss the medical evaluation and use of drop-testing experiments in the next section.

3. Medical evaluation

In this section we introduce relevant medical description conventions, describe the drop-testing experiments that were carried out, and then evaluate the observed injury on a medical basis. We carried out three phases of medical observation:

1. immediate observation after the drop-test;
2. macroscopic patho-anatomical analysis;
3. microscopic patho-anatomical analysis.

We describe all of them in sufficient detail. However, for better understanding of the observation results, the drop-testing setup is presented first.

3.1. Drop-testing experiments

A test setup based on the free-fall gravity principle was chosen for analysis, see Figure 9. The experimental setup consists of a vertical slide that is moved along two columns (motor-1), which positions the impactor to a desired drop height. Furthermore, a horizontal table is driven by a second motor (motor-2). With the latter it is possible to bring the




Fig. 9. Drop-testing setup for soft-tissue analysis. Motor 1 drives the sled to a desired drop height (equivalent to a desired impact speed) and releases the sled. In order to decouple the sled during the impact from the impactor, this is mounted in a low-friction housing that enables free motion along the impact direction. Motor 2 moves the table so that a series of impacts can be executed autonomously by moving the subject after each test such that an intact part can be hit next. With the setup we are able to measure velocity, acceleration, contact force, and contact pressure.

target to a desired lateral location. The connection between motor-1 and the slide is made by a crossbar that connects to the slide with two magnets and a security bolt. As soon as the desired height is reached and the test subject is moved into place, both magnets and the bolt are disabled. This causes the slide and the impactor to fall down on the subject. The impactor is fixed to a cylindrical shaft that is inserted with loose fit into a concentric housing on the slide. This housing is covered with a low-friction layer that permits the shaft to move easily in vertical direction during the impact. Therefore, the mass of the impactor is decoupled from the slide, allowing adjusting the impact mass between relatively low and large values. The remaining kinetic energy associated with the mass of the slide after the impact is absorbed by two passive dampers (this process does not interfere with the collision incident).

In order to analyze the general effects of contact geometry on injury with this setup, generic impactor structures that cover all relevant curvatures in robotics have to be found. However, selecting this set of primitives is non-trivial. In order to structure the problem, we first subdivide every robot that executes manipulation tasks into three parts to be covered:

1. robot structure;
2. end-effector;
3. tools and grasped objects.

Specialized tools or grasped objects such as knives, or other very complex tools cannot be covered with a finite set of impactors. Their treatment needs special care and presumably separate testing similar to our work in Haddadin et al. (2011). However, the robot structure, the end-effector, and a




Fig. 10. Set of basic 3D impact primitives. From left to right: spherical (radius r), cuboid (edge lengths a, b), pyramid edge (tip angle α), and edge (opening angle α , length l).

significant amount of tools/grasped objects can be covered with the following four primitives, see Figure 10:

1. P1: spherical impactor with parameter vector radius $\mathbf{a}_s = r$;
2. P2: cuboid with parameter vector edge lengths $\mathbf{a}_c = (a, b)$;
3. P3: pyramid edge with parameter vector symmetric tip angle $\mathbf{a}_p = \alpha$;
4. P4: edge with parameter vector opening angle and length $\mathbf{a}_e = (\alpha, l)$.

Finding the exactly required granularity $\Delta \mathbf{a}_i$ of the primitive parameters needs larger experimental investigations than we can provide here. Therefore, we are currently working on closing this gap and intend to initiate an international research effort on generating the according data. Nonetheless, from our experimental experience and the known large variation properties of biological tissue, it seems reasonable to cover following intervals:

1. $\mathbf{a}_s \in [1, 5 : 5 : 20]$ mm
2. $\mathbf{a}_{a,b} \in [5 : 5 : 30] \times [5 : 5 : 30]$ mm
3. $\mathbf{a}_p \in [10 : 10 : 120]^\circ$
4. $\mathbf{a}_e \in [10 : 10 : 120]^\circ \times [10 : 10 : 200]$ mm




Fig. 11. Selection of impact primitives for soft-tissue testing. The circled ones were investigated experimentally.

Of course, from a practical point of view one might think of additional primitives such as a pyramid edge with $\alpha = 90^\circ$ and a phase angle β being the varied parameter. However, from a conservativeness point of view P3 covers such structures already. Therefore, these primitives can be considered in future work.

In this paper we experimentally investigate three concrete primitives, see Figure 11. These three samples are labeled small sphere, large sphere, and wedge:

- wedge: $\alpha = 45^\circ$ with fillet radius $r = 0.2$ mm, width $l = 200$ mm and weight 2.7 kg;
- small sphere: radius $r = 5$ mm and weight 2.1 kg;
- large sphere: radius $r = 12.5$ mm and weight 2.2 kg.

These contact geometries cover already some typical industrial grippers, or parts of objects to be grasped. Tests were carried out at 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, and 4.0 m/s. The impactor masses were $\approx 2, 4, 6, 8, 10$ kg (depending on the particular impactor primitive). The setup is equipped with a contact force sensor and two accelerometers, measuring the sled and impactor acceleration, which are used for obtaining impact velocity. Furthermore, we are able to measure maximum stress with a pressure foil.

In the next subsection, we describe our evaluation methodology that was used for medical assessment of the observed injury.

3.2. Methodology

3.2.1. AO classification and macroscopic analysis In medical terms, the description of low-severity injury is usually considered as a secondary injury that accompanies, e.g., fractures. In this sense the AO classification of the ‘Arbeitsgemeinschaft für Osteosynthesefragen’ (AO) (Rüedi et al., 2007)⁶ is one of the most important international ones. The AO is also termed ASIF (Association for the Study of Internal Fixation) and aims to give a worldwide and definite fracture description of the human skeleton. Furthermore, there exists a subgroup within the AO classification, dealing with concomitant injuries, i.e. skin and soft-tissue injuries (muscle and tendon, nerve, and vessel). We

chose this sub-classification for precisely and objectively describing skin and soft-tissue injuries observed during the experimental series. It is important to note that this categorization enables others to comprehend and reproduce our results. The AO classification consists of a five-digit alphanumeric code for localization and description of fractures. Further codes for concomitant injuries in traumatology describe soft-tissue injuries.⁷ Each injury is associated with an alpha-numerical combination, for which the numbers range from 1 to 5, with injury severity generally increasing. Characters are assigned to skin injuries either as IC (injury closed) or IO (injury open). For muscle and tendon injuries they are denoted with MT for muscle and tendon and NV for nerve and vessel injuries. The purpose of the AO classification is to enhance communication between physicians, documentation, and research, leading to therapy choices found on properly assembled, clearly expressed, and readily accessible data, which are foundations for today’s evidence-based medicine. The parts of the AO classification that are relevant for our work, are

1. skin damage (I);
2. muscle and tendon injury (MT); and
3. nerve and vessel injury (NV).

Furthermore, these classes are detailed as follows.

Closed skin injury:

- IC1: no skin injury;
- IC2: contusion without skin opening;
- IC3: circumscribed décollement (avulsion);
- IC4: extensive, closed décollement (avulsion);
- (IC5: necrosis by deep contusion).

Open skin injury:

- (IO1: skin puncture from inside to outside);
- IO2: skin puncture from outside < 5 cm with contused margins’
- IO3: skin lesion > 5 cm, circumscribed décollement (avulsion) with marginal contusions;
- IO4: skin loss, deep contusion, abrasion;
- (IO5: extensive décollement (avulsion)).

Muscle and tendon injury

- MT1: no injury;
- MT2: circumscribed muscle injury (limited to one compartment);
- MT3: extensive muscle involvement (in two or more compartments);
- (MT4: avulsion or loss of a entire muscle groups, severed tendon);
- (MT5: compartment syndrome, crush syndrome).

Neurovascular injuries:

- NV1: no injury;
- NV2: isolated nerve lesion;
- NV3: circumscribed vascular injury;
- NV4: combined neurovascular injury;
- (NV5: subtotal- or total amputation).

In this paper we consider IC2 to be the appropriate threshold and label its occurrence ‘key impact’. In addition to the AO description, we measured the width, length, and depth of occurring wounds with a manual caliper. Prior to and after each impact series, images for documentation purposes were taken of the specimen. After the impacts were carried out, an initial observant description was recorded. After completing an impact series, the specimen was removed from the drop-testing device for profound investigation. First, the skin surface was evaluated and classified as IC1–IC5 or IO2–IO5. If no obvious skin opening could be observed, cuboid 1 cm³ skin samples were taken and fixed in formalin for microscopic evaluation. Then, the specimen was dissected layer by layer and injuries of the underlying fat tissue, muscles, fasciae, and serous membranes (peritoneum, pleura) were recorded. We slightly adjusted the AO classification as some outcomes cannot occur or be observed in our test series (the excluded parts are shown in brackets in the previous definition). Necroses by deep contusion, compartment syndrome, and crush syndrome (Rüedi et al., 2007) were removed, as these can only occur in living tissue due to being functional sequelae of tissue damage. Further avulsion or loss of an entire muscle group, severed tendon, skin puncture from inside to outside, extensive avulsion, and subtotal or total amputation cannot be observed in our test series. Since we are currently planning further tests using extremity parts, this may be investigated in detail in the future. However, the use of abdominal wall was chosen as experimental tissue, as it is comparatively easy to execute drop tests with equivalent impact conditions. The different successive layers have similar dimensions and, most importantly, they remain the same for each specimen. Using, for example, extremity specimens, these conditions are not met as the distance of the centrally lying bone to basically every surface position varies. This would only complicate the analysis and reduce readability.

Relevant definitions from the AO classification The following definitions can be found in Tutsch (2009) (some of them being similar for the ones in the review in Section 2). Contusion is a blunt organ injury (in our case skin injury) with the visible, compared with commotion, pooling of blood within the tissue, based on the rupture of small vessels (venules, capillaries, and arterioles). Avulsion (décollement) is the shearing or tearing off of one organ from another (e.g. tearing off the skin from the underlying subcutaneous tissues). However, it does not necessarily

mean open injury, which is termed open avulsion. Skin puncture is the penetration of dermal tissues down to the subcutaneous tissue, which does not necessarily mean the involvement of deeper lying tissues. We used the definition of circumscribed muscle injury as being muscle damage, involving one or two muscles layers. Extensive muscle involvement is the damage of three or, if present, more layers of the used abdominal wall specimen.

3.2.2. Microscopic analysis

Processing Microscopic analysis requires a multi-step processing of the investigated tissue (Lang, 2006). The essential steps are described hereafter.

The suspected area is excised in a cuboid manner of $\approx 1 \text{ cm}^3$ using a scalpel. The removed tissue sample is then placed in a container filled with 10% formalin solution and a buffer. The volume of the solution should be at least 10-fold to ensure the tissue sample is entirely soaked with formalin. The specimen is then stored for $\approx 24 \text{ h}$. Afterwards, the formalin-fixated tissue sample is manually cut into fine pieces. To dehydrate the sample, it is placed into an ascending array of alcohol solutions. The alcohol is then replaced by Xylol and phosphate-buffered saline for the final embedding in paraffin. The small paraffin (wax) block can then be placed in an automatic microtome in order to cut slices of $< 10 \mu\text{m}$. These thin slices are laid in a water bath to straighten them out and removed with a forceps onto a glass slide. On the glass slide, they are then stained with the commonly used hematoxylin–eosin stain. Staining requires the removal of paraffin by washing the slide again with Xylol and a descending array of alcohol. Hematoxylin stains all basophilic (acidic) structures such as nuclear DNA blue, whereas eosin stains all acidophilic (basic) structures, e.g. proteins, red.

Evaluation methodology The main goal of the microscopic analysis was to find a more detailed distinction between intact and injured skin, which might not be possible via macroscopic visual inspection only. It is important to note that the microscopic analysis is rather a tool of injury exclusion than of injury verification, as artifacts (excision damage, processing damage, etc.) may occur. However, the sensitivity of the method is unquestionable. Microscopic analysis begins with a scanning view over the entire sample at low magnification to unveil obvious processing mistakes (e.g. folded sample, unstained regions). Then, the scope is magnified to investigate the significant areas, i.e. the dermis and the subcutaneous tissue. Obvious dermal or subdermal defects are focused on and documented. The defects are divided into incision-, hollow-, groove-, notch-, rupture-, or stellate-like in shape and longitudinal, transverse, or oblique in direction. Negative results are termed ‘The dermis is fully intact. No definite defect visible’, see Figures 12 and 13.




Fig. 12. Impact area at 10× magnification (hematoxylin–eosin (HE)). Dermis is fully intact. No definite defect visible.




Fig. 13. Microscopic evaluation for wedge 45°, 4.7 kg, 1.5 m/s: 25× (left) and 40× (right) magnification (hematoxylin–eosin (HE)). In the corium there is an oblique running, relatively sharp rupture-like area. This could have been induced by the histological cutting conditions. No definite defect is visible.

3.3. Results

The presented evaluation is rather a summary of the observations that are useful for the design of the controller in Section 4 than an interpretative evaluation of the results. Figure 14 summarizes the risk curves that can be assembled from the acquired data. Note that the sole observation of key impacts alone is not sufficient if there is an entirely intact skin but also deeper tissue injury. This is especially important if this may involve nerves or arteries. This third group of soft-tissue injuries, nerves and vessels, is always feasible in the case of penetrative muscle injuries, as larger neurovascular structures lie beneath. Consequently, the actual key impact was chosen according to the expected resulting injury in a human being. This must be in the range of total reversibility (*restitutio ad integrum*), meaning that it should not leave any permanent damage. In the case of doubtful results the key impact was always chosen to be the more conservative interpretation. These impact results were then integrated into the real-time robot control as part of the injury database, see Section 4.

In Figure 14 (column one to three) the weight–velocity–injury scale relations are depicted for 276 impact tests. Each impactor (wedge, small sphere, and large sphere) occupies one column, showing the induced skin, muscle and tendon, and nerve and vessel injuries. The tissue damage severity (ranging from one to five in terms of AO) is equivalent to the grayscale rectangular fields. White fields delineate impacts that did not show tissue damage (IC1, MT1,

or NV1). Black areas reflect the highest severity possible. Please note that in ‘IC/IO for small sphere’ closed and open skin injuries are shown together in one graph. More specifically, the black rectangles denote open skin injuries and not closed injuries of IC5 (necrosis by deep contusion).

Skin damage caused by the wedge impactor is limited to contusions and closed décollement. Only weights above 8 kg and impact velocities of at least 3.0 m/s result in larger contusions and décollement. The most severe injury for the large sphere impactor are small avulsions of $\approx 10 \text{ mm}^2$. At velocities $< 2 \text{ m/s}$ the large sphere appears to induce only mild skin damages. Neither the larger sphere nor the wedge perforate the skin, which was also verified by microscopic analysis. On the other hand, the small sphere already pierces the skin at rather low velocities and weights. Therefore, safety can only be guaranteed with velocities $< 2 \text{ m/s}$ and $< 6 \text{ kg}$. The wedge impacts left negligible muscle damage in all impacts up to 1.5 m/s and even most injuries above can be considered as benign. The wedge shows a relatively large safety range for muscular tissue. Large sphere impactor muscle injuries with $> 8 \text{ kg}$ can be considered endurable up to 1.5 m/s. Results of the small sphere on the skin show similar severity on the muscle layer. The maximum velocity should not exceed 1.5 m/s and with additional weights the velocity must be reduced below 1.0 m/s. As mentioned above, the neurovascular injuries were chosen according to muscle penetration. However, they are not elaborated in this work as a reasonable assessment cannot be performed on non-vital tissue.

Following the initial tests with the three basic geometries, we conducted a series with a tool used in everyday industrial life, a Torx T45 bitset, see Figure 16. Figure 15 shows the observed injury for the different mass–velocity pairs. The hexalobular internal driving feature or Torx showed initially a similar severe injury pattern as the small sphere impactor. However, upon more detailed investigation, the open skin injuries appear later, more specifically at velocities of 4.0, 3.5, and 3.0 m/s at 4, 6, and 8 kg impact mass, respectively. Cutaneous defects at $< 1.0 \text{ m/s}$ and impact mass $< 8 \text{ kg}$ could not be observed. Muscular involvement occurred at $> 3.0 \text{ m/s}$ for $m = 2 \text{ kg}$. At an impactor weight of $> 6 \text{ kg}$, muscle damage can be definitely excluded only at velocities $< 0.5 \text{ m/s}$. Ponderable neurovascular involvement should only be expected at high weights and velocities, which could be explained by the length of the torx.

It should be noted that all experiments conducted up to this point imply certain conditions that deviate from real robot–human impacts. These are mainly associated with the use of non-vital tissue. Compared with a living subject, non-vital tissue lacks numerous characteristics, such as muscle tension, skin pre-stretch, and the possibility of reacting or evading. Furthermore, it is not possible to investigate particular functional damages that may occur, such as arterial/venous hemorrhage, pain, or neurologic deficits. Nevertheless, the conducted tests and their results reflect rather a worse-case scenario than understating possible outcomes.

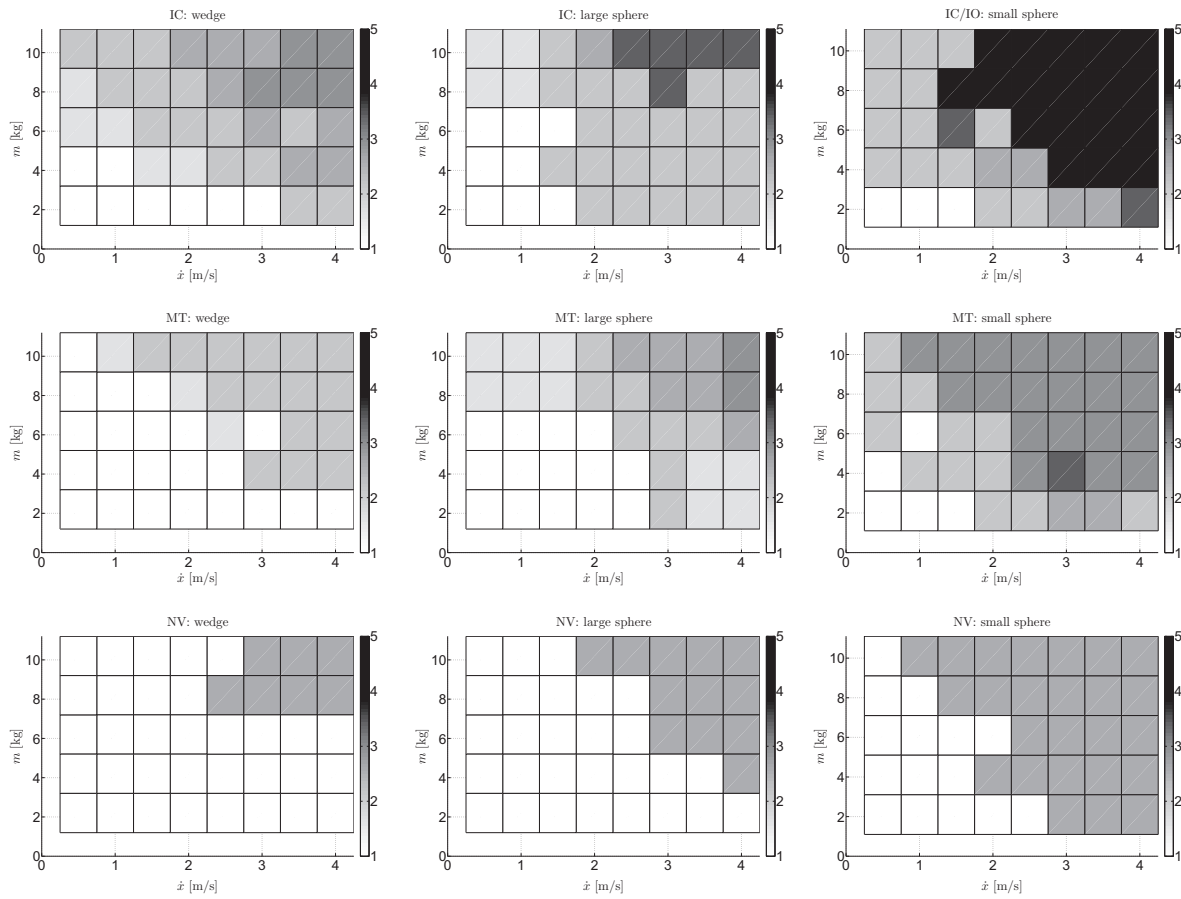


Fig. 14. Mass–velocity dependency of AO classification for the wedge (first column), large sphere (second column), and small sphere (third column). The upper row depicts the results for closed skin injury, the middle row for muscle and tendon injury, and the lower row for neurovascular injury.




Fig. 15. Mass–velocity dependency of AO classification for the Torx. The left plot depicts the results for closed skin injury, the center plot for muscle and tendon injury, and the right plot for neurovascular injury.



Fig. 16. Torx 45: a typical industrial end-effector tool.

In the next section we derive safety curves, which are suitable for real-time evaluation. They enable a robot to judge its current state (reflected inertia, instantaneous velocity, and curvature) regarding its potential safety characteristics in the case of an unexpected impact. We show how this representation can be used in closed-loop real-time fashion such that the robot does not exceed the respective medical limit values.

4. Knowledge-based real-time robot control

4.1. Safety curves

The intention of the analysis presented above is to understand how soft-tissue reacts under varying impact conditions. Furthermore, we seek the appropriate model parameters that are able to predict the occurrence of a particular injury for use in a safe motion controller. Despite the fact that we believe more experimental data is needed to fully understand the mapping $(mass, velocity, geometry) \rightarrow injury$ in general, such a full-range characterization of soft-tissue injury is certainly not necessary for robotic applications.

The prediction of speed limits for very low mass (< 1 kg) is of less interest mainly for two reasons:

- **Lower limit:** The reflected inertia of typical robots for interactive tasks is usually significantly larger (especially when being equipped with a gripper/hand and/or tools).
- **Upper limit:** A robot working side by side and/or cooperatively with a human should certainly not move faster than 4–5 m/s.⁸

Also speed limit evaluation for very high masses (> 20 kg) is not of large interest for service robotics, as:

- robots that are supposed to safely interact with humans are lightweight, leading to reflected inertias⁹ in the range of 1–15 kg; and
- even if the reflected inertia rises to high values (for example, near singularities) it does not make sense to

reduce speed under a certain value, or even stop the motion.

In the end, fixing a maximum speed limit (e.g. 4.5 m/s) for low inertia movements and a minimum speed limit (e.g. 0.1 m/s) for avoiding stopping of the robot close to singularities seems reasonable. Thus, all information needed can be obtained from our test results. The last unresolved question is which quantity reflects the correct injury threshold of the mass–velocity relation. Typically, one would aim for a description in terms of a physically interpretable quantity such as kinetic energy, contact force, or momentum. However, biomechanics research already showed that these do not necessarily capture human injury mechanics sufficiently. Therefore, so-called severity indices were developed for each body part. These map certain measurements such as acceleration, deflection, or force into an empirical criterion (e.g. HIC), which is then correlated to a certain injury type. This approach enables to equip anthropometric dummy devices with the respective sensors, evaluate severity indices in crash-testing setups, and map the recordings into injury scales such as the Abbreviated Injury Scale (AIS).¹⁰ Also in the early safety-related work in robotics, the HIC was applied as an intermediate measure of injury, which then needs to be mapped into an injury scale (Bicchi and Tonietti, 2004; Zinn et al., 2004). In other words, following relation is typically deduced: $(mass, velocity, curvature) \rightarrow (force, impact\ energy, HIC, \dots) \rightarrow injury$. However, due to the medical evaluation in terms of the AO classification in our work, the injury prediction problem needs no physical/severity index model, but is rather data driven, as it corresponds directly to the mapping $(mass, velocity, curvature) \rightarrow injury$. In particular, this relation is obtained from direct medical observation rather than the detor via an intermediate quantity. Owing to the complexity of the human injury mechanism, more consistent results can be presently expected in comparison to simplified model driven approaches that require validation steps and potentially suffer from larger inaccuracies. In the future, more complex and detailed modeling of the human injury biomechanics for further understanding of the problem is of course desirable and valuable. Presently, however, one could say that measurements of any physical quantity during an impact can be regarded as supplementary information only, but is not needed for describing the mapping $(mass, velocity, curvature) \rightarrow injury$ (of a certain body part). From associating velocity and mass values with ‘key impacts’, the resulting safety curves for the abdomen of the considered experiments are simply three limited regression lines in the $(mass, velocity)$ plane for a given impact primitive. The maximum velocity can be expressed as

$$v_{\max}(m) = reg.lim [c_1(i, \mathbf{a}_i)m + c_2(i, \mathbf{a}_i), v_1, v_2], \quad (1)$$

with $c_1(i, \mathbf{a}_i) < 0$ and $c_2(i, \mathbf{a}_i)$ being the coefficients of the safety curves for primitive i with primitive

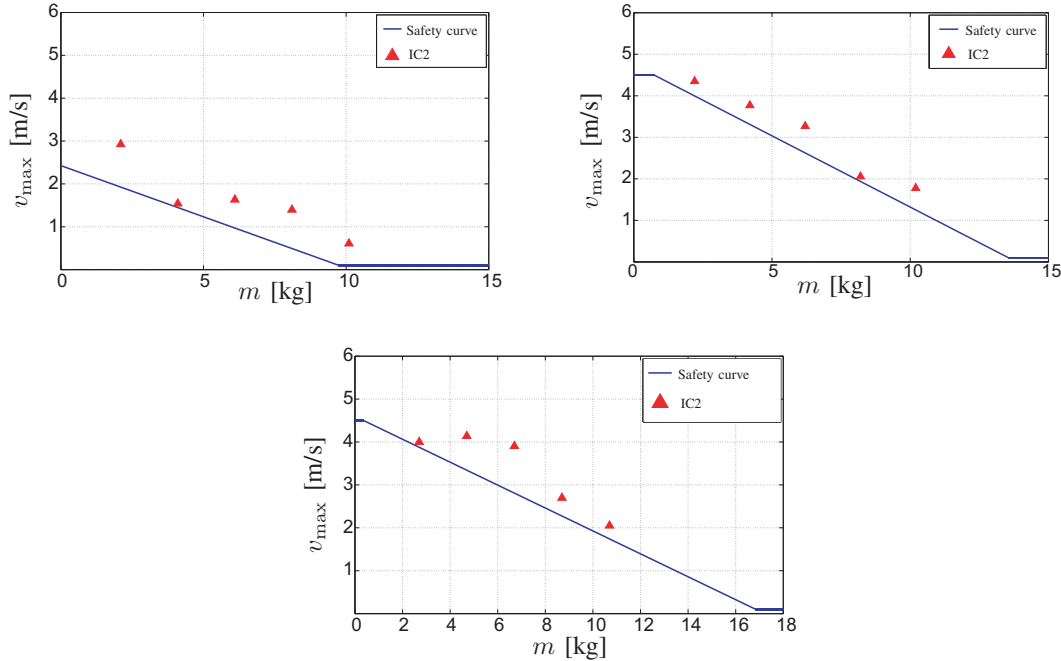


Fig. 17. Conservative safety curves for small sphere (top left), large sphere (top right), and wedge 45° (bottom) for robot real-time control. In order to remain well below all observed IC2 injuries, we chose simple shifted regression lines to represent the conservative injury estimates. If the conservative regression line does not already intersect with the y -axis before, the minimum and maximum velocity are chosen to be at 0.1 and 4.5 m/s, respectively.

parameters \mathbf{a}_i . The parameters v_1, v_2 denote the minimum and maximum velocity for the cutoff. Again, for isolation of the safety curves for robot motion control, no force sensor data is needed. Only the information at which velocity medically observable injury occurs at varying impactor masses is required. Sensor data, on the other hand, can be used for particular applications that need these quantities explicitly, such as force controlled tasks for prediction of contact forces.

Summarizing, the proposed methodology to obtain safety curves for robot speed limit control, is as follows:

1. select the geometric primitive to investigate;
2. carry out impact tests on pig subjects at varying contact speeds and impactor masses;
3. evaluate injury occurrence according to the AO classification and isolate ‘key impacts’;
4. data fitting of velocity and mass values associated with ‘key impacts’ (in the presented experiments simple bounded regression is sufficiently accurate);
5. delimiting the curve by maximum and minimum speed limits.

Figure 17 depicts the safety curves for the three tested primitives on the abdominal area. Maximum velocities are evaluated in the range of 0.1–4.5 m/s. The resulting lines are then shifted conservatively such that all data points are above the threshold. These curves build the basis to provide interpretable knowledge of human injury into real controllers of a robot, as described hereafter.

4.2. Injury database

The described impact-test results provide safety curves that relate maximum velocity and mass to injury for a distinct primitive (and body part), see Figure 17. Our goal is to enable a robot to use this knowledge for limiting the speed of moving parts such that an accidental collision would not cause injury above a given threshold in terms of the AO classification. In order to perform motions that are as fast as possible under the given safety constraint, the mapping $(mass, velocity, curvature) \rightarrow observed\ injury$ needs to be made accessible online. Such an injury database stores the coefficients of the safety curves (1) for each known primitive, so they can be used for real-time velocity shaping. In practice, end-effectors or relevant robot structures will not be uniquely described by a single primitive. They are rather complex geometric objects with varying characteristics. Owing to this diversity it is of course impractical to treat each robot/end-effector separately. In order to develop a generic approach, we formally decompose end-effectors into a set of rigidly coupled primitive objects. The overall geometric structure of the composed end-effector is described in terms of relative transformation matrices ${}^{EE}T_{Obj}$ between the end-effector and primitive object reference frame. Their respective geometric, dynamic (originally derived from CAD, dynamic identification, learning, etc.), and safety properties then adequately cover the robot hull.¹¹ Each relevant point of such a primitive object that shall be monitored is defined as a point of interest (POI). Each primitive object consists of a set of POIs, the position of its

center of mass $^{Obj}x_{COG}$, its mass m and inertia tensor I , and its relative pose with respect to the end-effector reference frame $^{EE}T_{Obj}$.

The following properties are associated with each *POI*. First, its relative pose with respect to the primitive object reference frame $^{Obj}T_{POI}$. Second, a set of geometric parameters *PARAMS* that represent the surface primitive *SURFACE*. Third, an identifier *SC-TYPE* for the type of safety curve (in our case limited regression). Fourth, a set of coefficients *COEFF* that describe the associated safety curves. Altogether, *COEFF*, *PARAMS*, and *SC-TYPE* compose the *PRIMITIVE* structure for every *POI*.

The structure of the resulting database is formally represented as follows:

$$\begin{aligned}
SoEEs &= \{SoObjects^k \times \{^{EE}T_{Obj}\}^k\} \\
^{EE}T_{Obj} &\in SE(3) \\
SoObjects &= \{POI^m \times \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbb{R}^{3 \times 3}\} \quad (2) \\
POI &= \{^{Obj}T_{POI} \times PRIMITIVE\} \\
^{Obj}T_{POI} &\in SE(3) \\
PRIMITIVE &= \{COEFF \times PARAMS \times SC-TYPE\} \\
COEFF &\in SoC(SC-TYPE) \\
PARAMS &\in SoP(SURFACE)
\end{aligned}$$

SoEEs is the set of end-effectors and *SoObjects* is the set of primitive objects.

Let us consider the case where a single large sphere primitive forms an object and end-effector. In other words, we assume the end-effector consists of one *POI* only that is associated with the large sphere. Overall, the database entry is as follows:

$$\begin{aligned}
SoEEs &= (SoObjects, ^{EE}T_{Obj}) \\
SoObjects &= (POI, [0 \ 0 \ 0.13]^T \text{ m}, 0.128 \text{ kg}) \\
^{EE}T_{Obj} &= \begin{pmatrix} I & [0 \ 0 \ 0.08]^T \text{ m} \\ 0 \ 0 \ 0 & 1 \end{pmatrix} \\
POI_1 &= (^{Obj}T_{POI}, SPHERE) \\
^{Obj}T_{POI} &= \begin{pmatrix} I & [0 \ 0 \ 0.1]^T \text{ m} \\ 0 \ 0 \ 0 & 1 \end{pmatrix} \quad (3) \\
SPHERE &= (COEFF, PARAMS, reg.lim) \\
COEFF &= (v_{\max}^{lim} = [0.1 \frac{\text{m}}{\text{s}}, 4.5 \frac{\text{m}}{\text{s}}] \\
&\quad [c_1 \ c_2] = [-0.3431 \frac{\text{m}}{\text{kg s}}, 4.7 \frac{\text{m}}{\text{s}}]) \\
PARAMS &= (r = 12.5 \text{ mm}).
\end{aligned}$$

Since this end-effector consists of a single primitive, the set of objects contains a single element only, as does the set of end-effectors. Of course, a real end-effector such as a gripper is composed of several primitives (e.g. gripper jaws), which in turn form the end-effector (full gripper).

As the injury database contains only injury relations for scalar masses and velocities, we need to obtain the instantaneous reflected mass of a given *POI* and its desired velocity for a given robot motion control command. Then, we can use the stored information such that the desired velocity can be scaled to remain below (1) in terms of potential injury.

4.3. Real-time dynamics based velocity shaping

4.3.1. *Reflected mass at POI* The well-known joint space dynamics formulation of a rigid robot is given by

$$M(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \boldsymbol{\tau}, \quad (4)$$

where $\mathbf{q} \in \mathbb{R}^n$ is the vector of joint positions, $M(\mathbf{q}) \in \mathbb{R}^{n \times n}$ is the inertia matrix, $C(\mathbf{q}, \dot{\mathbf{q}})$ is the centripetal and Coriolis matrix, $\mathbf{g}(\mathbf{q})$ the gravity vector, and $\boldsymbol{\tau}$ the joint torque. In joint coordinates the kinetic energy is

$$T_K = \frac{1}{2} \dot{\mathbf{q}}^T M(\mathbf{q}) \dot{\mathbf{q}}, \quad (5)$$

which can also be expressed in operational coordinates \mathbf{x} as

$$T_K = \frac{1}{2} \dot{\mathbf{x}}^T \Lambda(\mathbf{x}) \dot{\mathbf{x}}, \quad (6)$$

where $\Lambda(\mathbf{x})$ is the kinetic energy matrix associated with the operational space (Khatib, 1995). The relation between joint and operational space velocities $\dot{\mathbf{x}} = J(\mathbf{q})\dot{\mathbf{q}}$ is given via the according Jacobian $J(\mathbf{q}) \in \mathbb{R}^{6 \times n}$. The relation between $M(\mathbf{q})$ and $\Lambda(\mathbf{x})$ was also derived in Khatib (1995) and is well known to be

$$\Lambda(\mathbf{x}) = (J^T)^{-1}(\mathbf{q}) M(\mathbf{q}) J^{-1}(\mathbf{q}) \quad (7)$$

for non-redundant robots and

$$\Lambda(\mathbf{x}) = (J(\mathbf{q}) M(\mathbf{q})^{-1} J^T(\mathbf{q}))^{-1} \quad (8)$$

for redundant robots. Based on a decomposition of the inverse of the kinetic energy matrix into

$$\Lambda^{-1}(\mathbf{q}) = \begin{bmatrix} \Lambda_v^{-1}(\mathbf{q}) & \bar{\Lambda}_{v\omega}(\mathbf{q}) \\ \bar{\Lambda}_{v\omega}^T(\mathbf{q}) & \Lambda_\omega^{-1}(\mathbf{q}) \end{bmatrix} \quad (9)$$

and rewriting the Jacobian as

$$J(\mathbf{q}) = \begin{bmatrix} J_v(\mathbf{q}) \\ J_\omega(\mathbf{q}) \end{bmatrix}, \quad (10)$$

we may obtain a scalar value that represents the mass perceived at the end-effector, given a force in unit direction \mathbf{u} (Khatib, 1995). This quantity is denoted as the reflected robot inertia m_u in the \mathbf{u} direction. The derivation of the reflected rotational robot inertia I_u , perceived at the *POI* about direction \mathbf{u} is done analogously:

$$\begin{aligned}
m_u &= [\mathbf{u}^T \Lambda_v^{-1}(\mathbf{q}) \mathbf{u}]^{-1}, \quad (11) \\
I_u &= [\mathbf{u}^T \Lambda_\omega^{-1}(\mathbf{q}) \mathbf{u}]^{-1}.
\end{aligned}$$

In accordance with the drop testing outcome, the quantity m_u is needed for deriving the maximum allowable operational point velocity in \mathbf{u} direction that satisfies the safety curves, see Figure 17.

Next, we derive how to obtain the maximum allowable velocity of an operational space point, which is the basis for obtaining a safe maximum robot velocity.

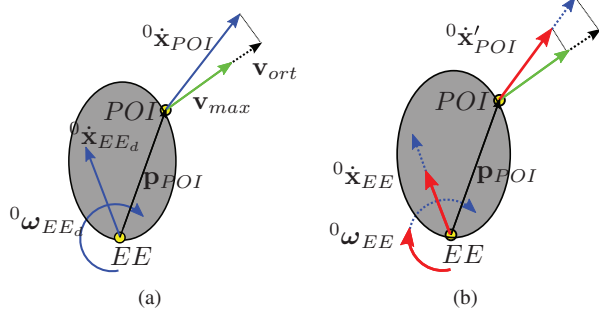


Fig. 18. Calculation of the safe POI velocity in the SMU based on the instantaneous reflected inertias, velocities, and associated curvature primitives for each POI .

4.3.2. Injury-based velocity shaping For a generic POI following scheme leads to its safe velocity.¹² The robot base coordinate system is denoted $\{0\}$ and the end-effector frame (operational frame) as $\{EE\}$.

1. Evaluate the unit vector \mathbf{u} that is perpendicular to the object surface in the POI (z -axis direction in POI coordinates):

$$\mathbf{u} = {}^0R_{POI}\mathbf{z}_{POI} \quad (12)$$

where ${}^0R_{POI}$ is the rotation matrix from POI coordinates to $\{0\}$ coordinates.

2. Calculate ${}^0\mathbf{v}_{POI}$ from the desired end-effector velocity ${}^0\mathbf{v}_{EE_d} = [{}^0\dot{\mathbf{x}}_{EE_d} \quad {}^0\boldsymbol{\omega}_{EE_d}]^T$:

$${}^0\mathbf{v}_{POI} = \begin{bmatrix} {}^0\dot{\mathbf{x}}_{POI} \\ {}^0\boldsymbol{\omega}_{POI} \end{bmatrix} = \begin{bmatrix} I_3 & -\hat{\mathbf{p}}_{POI} \\ 0_3 & I_3 \end{bmatrix} \begin{bmatrix} {}^0\dot{\mathbf{x}}_{EE_d} \\ {}^0\boldsymbol{\omega}_{EE_d} \end{bmatrix}, \quad (13)$$

where $\mathbf{p}_{POI} = [p_{x,POI} \quad p_{y,POI} \quad p_{z,POI}]^T$ is the position vector of the POI from the $\{EE\}$ origin. The matrix $\hat{\mathbf{p}}_{POI}$ is

$$\hat{\mathbf{p}}_{POI} = \begin{bmatrix} 0 & -p_{z,POI} & p_{y,POI} \\ p_{z,POI} & 0 & -p_{x,POI} \\ -p_{y,POI} & p_{x,POI} & 0 \end{bmatrix}. \quad (14)$$

3. Evaluation of the inverse of the operational space kinetic energy matrix at POI :

$$\Lambda_{v,POI}^{-1} = J_{v,POI}M^{-1}J_{v,POI}^T. \quad (15)$$

4. Reflected mass in the \mathbf{u} direction via (11):

$$m_{POI} = 1/(\mathbf{u}^T \Lambda_{v,POI}^{-1} \mathbf{u}) \quad (16)$$

5. Evaluation of the maximum velocity v_{\max} for m_{POI} , using the safety functions:

$$v_{\max}(m) = \text{reg.lim}[c_1(i, \mathbf{a}_i)m + c_2(i, \mathbf{a}_i), v_1, v_2], \quad (17)$$

6. Comparison of v_{\max} and the projection of ${}^0\mathbf{v}_{POI}$ in \mathbf{u} direction \mathbf{v}_{ort} : if $\|\mathbf{v}_{ort}\| \leq \|v_{\max}\|$, the desired speed remains. If $\|\mathbf{v}_{ort}\| > \|v_{\max}\|$, the new velocity ${}^0\mathbf{v}'_{POI}$ is

$${}^0\mathbf{v}'_{POI} = {}^0\mathbf{v}_{POI} \frac{\|v_{\max}\|}{\|\mathbf{v}_{ort}\|}.$$

7. Finally, the new end-effector velocity ${}^0\mathbf{v}_{EE}$ is

$${}^0\mathbf{v}_{EE} = \begin{bmatrix} {}^0\dot{\mathbf{x}}_{EE} \\ {}^0\boldsymbol{\omega}_{EE} \end{bmatrix} = \begin{bmatrix} I_3 & \hat{\mathbf{p}}_{POI} \\ 0_3 & I_3 \end{bmatrix} \begin{bmatrix} {}^0\dot{\mathbf{x}}'_{POI} \\ {}^0\boldsymbol{\omega}'_{POI} \end{bmatrix}. \quad (18)$$

This procedure is repeated for every POI and the most conservative ${}^0\mathbf{v}_{EE}$ is selected to be ${}^0\mathbf{v}'_{EE}$.

Next, we discuss some experiments, showing the performance of the system in basic ‘Move from pose A to pose B’ tasks. For this, an articulated robot arm is equipped with an end-effector, consisting of the investigated primitives from the drop-testing experiments.

4.4. Experiment

Before going into the details of the experiments, let us quickly review our ‘human-friendly’ robot control framework *Beasty*,¹² into which the database and the velocity supervision are embedded.

4.4.1. System architecture The task control unit (TCU) and the robot control unit (RCU) are the main parts of our robot control system architecture *Beasty* (Parusel et al., 2011) for the LWR-III. They serve as the general interface to the robot and communicate with each other via asynchronous protocols. The TCU is the general state-based control entity, which runs in non-real-time and provides the nominal robot actions and behaviors to the RCU. The RCU in turn runs at the same clock rate as the robot, assigning control, motion generation, interaction, and safety methods, i.e. the concrete behaviors. Furthermore, it interprets and validates the behaviors provided by the TCU. The new safe motion unit (SMU) constitutes the biomechanics-based supervision entity that implements the algorithm from Section 4.3.2, complementing the basic interaction and safety control schemes running in the RCU. It is the central part of the real-time speed limit control and receives information from the TCU, RCU, and the injury database (INJURY DB). Figure 19 depicts a schematic description of how the injury database interacts with the SMU. Information of the end-effector (object composition) is sent from the TCU to the SMU. This information is then used to approximate the end-effector by stored primitives. Then, the SMU retrieves the related injury data associated with the set of objects from the Injury DB (POI s, primitives, and injury coefficients). The RCU provides the dynamics information of the manipulator ($\mathbf{q}, \dot{\mathbf{q}}, M(\mathbf{q}), {}^0J_{EE}(\mathbf{q})$) and the current desired motion of the end-effector to the SMU via a real-time protocol running at 1 kHz. The SMU uses this data to calculate the reflected inertia at each POI -direction, generates ${}^0\mathbf{v}_{EE}$, and sends this back to the RCU. In turn, the RCU shapes its commands to the robot accordingly.

4.4.2. Experimental results In order to show the effectiveness of the developed injury-based algorithm, we equipped an LWR-III with an end-effector that is composed of the considered primitives from the drop-testing experiments,




Fig. 19. Safety management unit (SMU) embedded into the DLR real-time control framework *Beasty*.




Fig. 20. End-effector coordinates expressed in the robot base frame for the experiment. The 0y axis points towards the small sphere primitive.




Fig. 21. End-effector composed of impactor primitives and the associated POIs. POI1 is attached to the small sphere, POI2, POI3 to the edge, and POI4 to the large sphere. Please note that the corners of the end-effector have not been drop-tested yet. Therefore, we omit their effect on the POI selection, i.e. we assume the corners to be roundish.




Fig. 22. Trajectories for the ‘line test’ (left) and the ‘ribbon test’ (right). The ‘line test’ test shows the effect on the two spherical POIs and the ‘ribbon test’ is used for showcasing the behavior during a motion that combines scaling on all POIs.

see Figures 20–21. For this end-effector we selected four POIs: two on the tip of the spheres (POI1 and POI4 in Figure 21) and two on the wedge (POI2 and POI3). For the wedge two POIs need to be chosen, as the width of the wedge has a significant effect during rotational motion. Considering, for instance, two POIs at the edges of the wedge, one of them is certainly the fastest point of the primitive. However, the experiments carried out so far cannot serve as an injury knowledge source for the corners of the wedge, as the analysis of corners is yet to be done. A compromise for the experiment is to assign the POIs to the distal ends of the primitive, where the speed difference compared with those of the corners is negligible.

Since the LWR-III is explicitly designed for sensitive and safe interaction, one of its most important characteristics is the lightweight design. Since with this tool the robot is not able to execute potentially ‘unsafe’ motions, given its maximum speed and inertial characteristics, we shifted the safety curves such that the effect of the SMU becomes visible also for this manipulator. For the experiment, we selected a scaling factor of 0.2, while for the Torx 45 we would need to scale by 0.3.

Experiments were carried out for two different motions. In the line test (Figure 22 left) the end-effector moves between two locations laterally along the 0y -direction. In the ribbon test (Figure 22 right) the robot moves across four different positions, combining vertical and horizontal motions.

Figure 23 depicts the result for the line test. The end-effector is supposed to move between the two configurations at 1.5 m/s desired velocity. The SMU, in turn, limits the speed, depending on the respective motion direction and reflected inertia. During motion towards the positive 0y -direction, the critical POI is POI1 and along negative direction POI4, respectively. Since the safety curve associated with POI1 is more restrictive than that for POI4 (the small sphere is more dangerous than the larger one) the SMU reduces the maximum speed stronger in the positive 0y -direction.

In Figure 24, the ribbon test results are shown. Since the operational frame ($\{0\}$ -coordinates) is rotated with respect to the Cartesian world frame ($\{W\}$ -coordinates), the

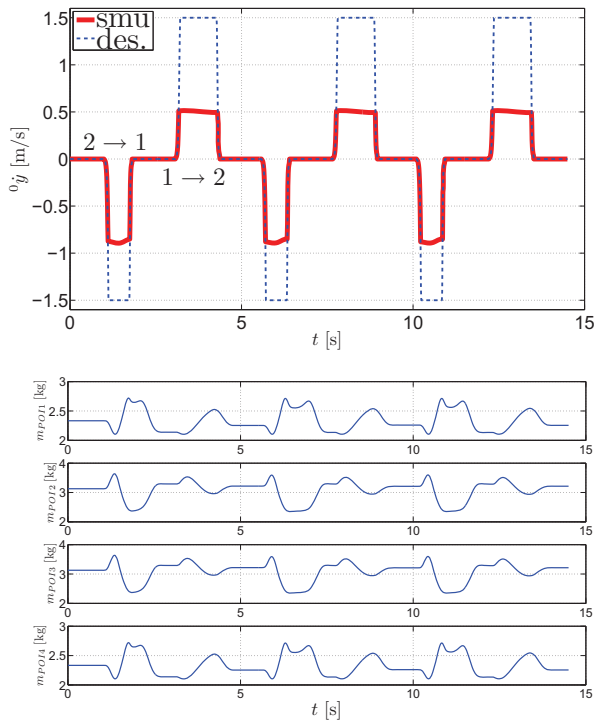


Fig. 23. SMU experiments with the ‘line test’. Cartesian motion along the Wx -axis (0y in the end-effector frame) and reflected inertia evaluated in each POI ’s direction.




Fig. 24. SMU experiments with the ‘ribbon test’. Cartesian motion along the Wz -axis (0x - and 0z -axis on the end-effector frame) and reflected inertia evaluated in each POI ’s direction.

velocity is depicted for all three dimensions. Lateral movements are represented in the 0y -direction. Vertical movements have components in both the 0x - and the 0z -direction. Accordingly, the operational frame, displacements 1–2 and 3–4 in Figure 22 (rising up movement of the end-effector) correspond to positive 0x and negative 0z in Figure 24. Motion segments 2–3 and 4–1, in contrast, represent negative 0x and positive 0z motion. For the lateral direction, displacements 2–3 and 4–1 correspond to negative and positive 0y motions. In Figure 24 one can observe that the SMU limits the 4–1 motion stronger than the 2–3 motion. This is again due to the fact that along this direction $POI1$ is relevant, which is associated with the small sphere. However, for this case $POI2$ and $POI3$ are also involved in the speed limitation. Owing to the symmetry of the system they equally contribute to the velocity limits along 4–1 and 2–3. Finally, it is possible to observe that speeds in rising up motions (1–2 and 3–4) are not yet limited by the controller. This is because POI s were placed only in frontal and lateral positions, and none on the back of the end-effector. The proposed algorithm runs on a standard PC under VxWorks at a rate of 1 kHz, which is typically an order of magnitude faster than typical industrial robot arm controllers.

4.5. Standardization

As already mentioned, current standardization efforts aim at incorporating findings such as those developed in the present work. In this context we believe that our results fit the according needs very well and could serve as a template on how to generate the required data sets. In particular, our approach allows the integration of other data and is not specific to our evaluation schemes. Of course, the acceptance of such an approach would need an international consensus and a consequential internationally coordinated research effort for completion and validation of data. This would ensure the cross-fertilization of results and a coordinated path towards generating sufficient data covering all relevant applications.

Apart from the basic data and unified representation we provide, it is important to implement the proposed algorithms according to relevant safety ratings. In particular, the robot dynamics, the guarantee of keeping a desired velocity, and the correct implementation of our safe velocity controller itself are prime. Together with a formal certification process chain that is based on a separation of responsibility it would be possible to safely implement the concept. Following roles could be assigned for this:

1. **robot manufacturer:** provides robot curvatures (e.g. CAD), inertial data, and safe velocity controller;
2. **mechanical designer:** assigns primitives to the robot structure, tools, and end-effectors;
3. **application designer/integrator:** analyzes possible worst cases in target application, i.e. risk assessment, therefore select the injury risk curves applicable to the specific application.

5. Conclusion

The contributions of this paper are fourfold. First, drop-testing experiments with pig abdominal wall samples at varying mass, velocities, and impactor geometries have been conducted. These were carried out to generate fundamental injury analysis data that is still missing even in biomechanics and forensics.

Second, we have proposed a procedure for using medical judgement and classification to generate appropriate representations of injury knowledge for further processing. This was done according to the AO classification. In general, the AO classification allows an objective labeling of (soft-tissue) injury based on medical observation. We have selected IC2 (contusion without skin opening) as a maximally tolerable threshold of injury.

Third, we have developed safety curves for representing ‘safe’ robot speed for a given configuration, mass, impact geometry, and impacted body part. We have also designed an online injury database architecture that makes the generated insights accessible in real-time.

Finally, we have designed and experimentally verified a real-time control scheme that limits the end-effector velocity according to the current reflected inertia and the geometric properties of the end-effector. This ensures that a possible collision with a human (abdomen) cannot generate injury above a certain injury level by exploiting the knowledge from the injury database. In addition, our previous work on blunt impact testing can be easily integrated into the architecture and algorithms, which makes the approach very generic.

Our future work will concentrate on the following aspects. First of all, we are already planning additional large-scale impact tests for increasing the amount of data. In particular, we want to understand injury mechanisms of other body parts. Also the extension to living tissue is an interesting direction, as the considered subjects show rather conservative injury responses (e.g. no muscle tension). From an algorithmic point of view the automatic identification of *POIs* and their relative primitives of a generic end-effector, starting from CAD models or even from a camera view is another direction we intend to pursue.

Finally, we believe our results are also valuable for advancing model-based techniques. In addition to deriving the relation (*mass, velocity, curvature*) \rightarrow *injury*, understanding injury biomechanics also involves an insight into the physical processes, which is particularly important for further generalization, e.g. in terms of scaling. However, it is clear that the respective biomechanical data needs to be derived first for setting up according models and validate them rigorously.

Notes

1. The Berufsgenossenschaft (BG) is the German employer’s liability insurance association.
2. Both studies analyzed spherical contacts and resulting blunt soft-tissue injury in terms of lacerations or contusions.

3. In the future, we extend the analysis to various other body parts and also to lateral impacts that may cause abrasions and cuts. However, this is an enormous amount of data that would somewhat cloak the essence of the work. Therefore, we intentionally treat a single body part, namely the abdominal area, which is due to its regular structure more clear to evaluate.
4. The intention of the present work is to develop a methodology rather than providing raw masses of data. The accumulation of this data for more primitives and body parts is currently in progress. However, it is certainly beyond the scope of the present paper to discuss and evaluate them in detail as well.
5. A total of 267 drop tests were conducted to generate the data for this paper. We are currently working on producing more test data in order to gather also the statistics of the tests. However, the results show extremely good consistency overall.
6. Parts of it are also known as the Müller classification, named after Maurice Edmond Müller. He was a Swiss surgeon and is known as a pioneer in orthopedic surgery.
7. Further descriptions exist as well.
8. In fact, as we showed in our previous work ≈ 2 m/s is a reasonable speed limit. This is derived from blunt head impacts.
9. Please note that we do not refer to singularities. Singularities have to be analyzed differently, as even though reflected inertia approaches infinity, velocity goes to zero at the same time, i.e. the kinetic energy is limited.
10. See <http://www.aqc.ch/scores.htm>.
11. Please note that we do not provide an algorithm to do this automatically. This is left for future work.
12. Please note that we omit all indices for sake of clarity. Furthermore, the dependency on \mathbf{q} is also omitted if obvious.
13. *Beasty: Beyond Industrial Safety*.
14. Presumably, this would be the case for any inelastic measurement device such as the well-known Tekscan pressure sensor, as well.

Funding

This work has been partially funded by the European Commission’s Sixth Framework Programme as part of the project SAPHARI (grant number 287513). We would also like to acknowledge the support of KUKA Roboter GmbH.

References

- Albu-Schäffer A, Haddadin S, Ott C, Stemmer A, Wimböck T and Hirzinger G (2007a) The DLR Lightweight Robot: Lightweight design and soft robotics control concepts for robots in human environments. *Industrial Robot Journal* 34: 376–385.
- Albu-Schäffer A, Ott C and Hirzinger G (2007b) A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *The International Journal of Robotics Research* 26: 23–39.
- BG/BGIA (2009) Empfehlungen zur Gestaltung von Arbeitsplätzen mit kollaborierenden Robotern. Technical report.
- Bicchi A and Tonietti G (2004) Fast and soft arm tactics: dealing with the safety-performance trade-off in robot arms design and control. *IEEE Robotics and Automation Magazine* 11: 22–33.
- Bischoff R, Kurth J, Schreiber G, et al. (2010) The KUKA-DLR lightweight robot arm: a new reference platform for robotics research and manufacturing. In *International Symposium on Robotics (ISR2010)*, Munich, Germany.
- Cainelli T, Giannetti A and Rebora A (2008) *Manuale di Dermatologia Medica e Chirurgia*, 4th edn. New York: McGraw-Hill.

- De Luca A and Mattone R (2004) An adapt-and-detect actuator FDI scheme for robot manipulators. *IEEE International Conference on Robotics and Automation (ICRA2004)*, New Orleans, LA, pp. 4975–4980.
- Edlich R, Rodeheaver G, Morgan R, Berman D and Thacker J (1988) Principles of emergency wound management. *Annals of Emergent Medicine* 17: 284–302.
- Eisenmenger W (2004) *Handbuch gerichtliche Medizin - Spitze, scharfe und halbscharfe Gewalt*. Berlin: Springer-Verlag (in German).
- Fung Y (1993) *Biomechanics: Mechanical Properties of Living Tissues*. New York: Springer.
- Gao D and Wampler C (2009) On the use of the head injury criterion (HIC) to assess the danger of robot impacts. *IEEE Robotics and Automation Magazine* 16(4): 71–74.
- Goldsmith P, Francis B and Goldenberg A (1999) Stability of hybrid position/force control applied to manipulators with flexible joints. *International Journal of Robotics and Automation* 14: 146–159.
- Haddadin S, Albu-Schäffer A, Frommberger M, Roßmann J and Hirzinger G (2009a) The ‘DLR Crash Report’: towards a standard crash-testing protocol for robot safety - part I: results. In *IEEE International Conference on Robotics and Automation (ICRA2008)*, Kobe, Japan, pp. 272–279.
- Haddadin S, Albu-Schäffer A and Hirzinger G (2007a) Safe physical human–robot interaction: measurements, analysis and new insights. In *International Symposium on Robotics Research (ISRR2007)*, Hiroshima, Japan, pp. 439–450.
- Haddadin S, Albu-Schäffer A and Hirzinger G (2007b) Safety evaluation of physical human–robot interaction via crash-testing. *Robotics: Science and Systems Conference (RSS2007)*, Atlanta, GA.
- Haddadin S, Albu-Schäffer A and Hirzinger G (2009b) Requirements for safe robots: measurements, analysis and new insights. *The International Journal of Robotics Research* 28: 1507–1527.
- Haddadin S, Albu-Schäffer A and Hirzinger G (2010) Soft-tissue injury in robotics. *IEEE International Conference on Robotics and Automation (ICRA2010)*, Anchorage, AK, pp. 3426–3433.
- Haddadin S, Albu-Schäffer A, Haddadin F, Rossman J and Hirzinger G (2011) Experimental safety study on soft-tissue injury in robotics. *IEEE Robotics and Automation Magazine* 18(4): 20–34.
- Haddadin S, Albu-Schäffer A, Luca AD and Hirzinger G (2008) Collision detection and reaction: A contribution to safe physical human–robot interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2008)*, Nice, France, pp. 3356–3363.
- Heinzmann J and Zelinsky A (2003) Quantitative safety guarantees for physical human–robot interaction. *The International Journal of Robotics Research* 22: 479–504.
- Hogan N (1985) Impedance control: An approach to manipulation: Part I - theory, Part II - implementation, Part III - applications. *Journal of Dynamic Systems, Measurement and Control* 107: 1–24.
- Ikuta K, Ishii H and Nokata M (2003) Safety evaluation method of design and control for human-care robots. *The International Journal of Robotics Research* 22: 281–298.
- International Organization for Standardization (2006) ISO10218: Robots for industrial environments – Safety requirements – Part 1: Robots.
- International Organization for Standardization (2011) ISO10218-1: Robots and robotic devices – Safety requirements for industrial robots – Part 1: Robots.
- Iwata H and Sugano S (2009) Design of human symbiotic robot TWENDY-ONE. *IEEE International Conference on Robotics and Automation (ICRA2009)*, Kobe, Japan, pp. 580–586.
- Khatib O (1995) Inertial properties in robotic manipulation: an object-level framework. *The International Journal Robotics Research* 14: 19–36.
- Kulic D and Croft E (2007) Pre-collision strategies for human robot interaction. *Autonomous Robots* 22: 149–164.
- Lang G (2006) *Histotechnik*. Berlin: Springer.
- Lim H-O and Tanie K (2000) Human safety mechanisms of human-friendly robots: Passive viscoelastic trunk and passively movable base. *The International Journal of Robotics Research* 19: 307–335.
- Macchiarelli L, Arbarello P, Cave Bondi C, Di Luca NM and Feola T (2002) *Medicina Legale (compendio) II edition*, Turin: Minerva Medica Publications.
- Matthias B, Oberer-Treitz S, Staab H, Schuller E and Peldschus S (2010) Injury risk quantification for industrial robots in collaborative operation with humans. *Robotics (ISR), 2010 41st International Symposium on and 2010 6th German Conference on Robotics (ROBOTIK)*, pp. 1–6.
- Meyer W (1996) Comments on the suitability of swine skin as a biological model for human skin (German). *Der Hautarzt* 47: 178–182.
- Moore KL, Dalley AF and Agur AMR (2009) *Clinically Oriented Anatomy*, 6th edition. Baltimore: Lippincott Williams & Wilkins.
- Oberer S and Schraft R-D (2007) Robot-dummy crash tests for robot safety assessment. In *IEEE International Conference on Robotics and Automation (ICRA2007)*, Rome, Italy, pp. 2934–2939.
- Ogrodnikova O (2009) How safe the human–robot coexistence is? *Acta Polytechnica Hungarica* 6(4): 51–74.
- Park J, Haddadin S, Song J and Albu-Schäffer A (2011) Designing optimally safe robot surface properties for minimizing the stress characteristics of human–robot collisions. In *IEEE International Conference on Robotics and Automation (ICRA2011)*, Shanghai, China, pp. 5413–5420.
- Park J-J and Song J-B (2009) Collision analysis and evaluation of collision safety for service robots working in human environments. In *International Conference on Advanced Robotics, 2009 (ICAR 2009)*, pp. 1–6.
- Parusel S, Haddadin S and Albu-Schäffer A (2011) Modular state-based behavior control for safe human–robot interaction: a lightweight control architecture for a lightweight robot. In *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4298–4305.
- Povse B, Koritnik D, Kamnik R, Bajd T and Munih M (2010) Industrial robot and human operator collision. In *IEEE International Conference on Systems Man and Cybernetics (SMC2010)*, Istanbul, Turkey, pp. 2663–2668.
- Ross MH and Pawlina W (2006) *Histology: A Text and Atlas With Correlated Cell and Molecular Histology*. Baltimore: Lippincott Williams & Wilkins.
- Rüedi TP, Buckley RE and Morgan CG (2007) *AO Principles of Fracture Management*. vol 1, 2nd edition. Stuttgart: Thieme Verlag.

- Shergold O and Fleck N (2004) Mechanisms of deep penetration of soft solids, with application to the injection and wounding of skin. *Proceedings of the Royal Society London A* 460: 3037–3058.
- Shergold OA and Fleck NA (2005) Experimental investigation into the deep penetration of soft solid by sharp and blunt punches, with application to the piercing of skin. *Journal of Biomechanical Engineering* 127: 838–848.
- Shergold OA, Fleck NA and Radford D (2006) The uniaxial stress versus strain response of pig skin and silicone rubber at low and high strain rates. *International Journal of Impact Experiments* 32(9): 1384–1402.
- Shin D, Sardellitti I and Khatib O (2008) Hybrid actuation approach for human-friendly robot design. In *IEEE International Conference on Robotics and Automation (ICRA2008)*, Pasadena, CA, pp. 1741–1746.
- Stemmer A, Albu-Schäffer A and Hirzinger G (2007) An analytical method for the planning of robust assembly tasks of complex shaped planar parts. In *International Conference on Robotics and Automation (ICRA2007)*, Rome, Italy, pp. 317–323.
- Townsend W and Salisbury J (1993) Mechanical design for whole-arm manipulation. In Dario P, Sandini G and Aebischer P (eds), *Robots and Biological Systems: Towards a New Bionics?* New York: Springer, pp. 153–164.
- Tutsch D (2009) *Roche LexikonMedizin*. Munich: Urban & Fischer Verlag.
- Versace J (1971) A review of the severity index. *Proceedings of the 15th Stapp Conference*, SAE Paper No.710881, pp. 771–796.
- Wassink M and Stramigioli S (2007) Towards a novel safety norm for domestic robots. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2007)*, San Diego, CA, pp. 3243–3250.
- Yamada Y, Hirasawa Y, Huand S and Umetani Y (1996) Fail-safe human/robot contact in the safety space. *IEEE International Workshop on Robot and Human Communication*, pp. 59–64.
- Yamada Y, Hirasawa Y, Huang S, Umetani Y and Suita K (1997) Human-robot contact in the safeguarding space. *IEEE/ASME Transactions on Mechatronics* 2: 230–236.
- Zinn M, Khatib O, Roth B and Salisbury J (2004) Playing it safe. *IEEE Robotics and Automation Magazine* 11(2): 12–21.
- Zinn M, Khatib O, Roth B and Salisbury JK (2005) A new actuation approach for human-friendly robot design. *International Journal of Robotics Research* 23: 379–398.
- Zollo L, Siciliano B, De Luca A, Guglielmelli E and Dario P (2005) Compliance control for an anthropomorphic robot with elastic joints: Theory and experiments. *ASME Journal of Dynamic Systems, Measurements and Control* 127: 321–328.

Appendix: Measurements and safety curves

This appendix provides supplementary experimental data from the drop-testing experiments and gives further insight into the understanding of soft-tissue injury.

A.1. Impact forces over time

In Figures 25–27 the force–time plots are depicted for the three impactors at varying velocity. For each impactor the

forces are shown for every impact mass. From this data some interesting observations can be drawn.

For the wedge one can see rather bell-shaped force–time curves, which can be interpreted as a simple spring contact in first approximation. The same observation holds for the spheres at low masses and low velocities. For the large sphere one observes a small drop in force at impacts with $m \geq 4.2$ kg and at higher speeds, see Figure 26. This indicates a rupture of the underlying muscle tissue (consistent with medical observations from Section 3.3) with the skin being still intact. For the small sphere significantly larger force drops are observed especially at high velocities and large impact masses, see Figure 27. For these cases, the impactor penetrated the skin layer with underlying muscle tissue.

To sum up, the force profiles comply in their general behavior with the medical observations, as muscle ruptures can be observed as discontinuities in the contact force over time. In particular, the small sphere causes this effect, which is presumably due to the larger impact energy density.

A.2. Impact forces versus impact velocity

Figure 28 depicts the maximum force as a function of contact speed and parameterized by the impactor mass for every primitive. Assuming a rather linear relation between force and penetration, it is evident that for a constant mass value, force increases linearly with velocity until penetration occurs. After the linear region, a change of trend can be observed, and a saturation effect is present for the small sphere. This is presumably due to the severe damage to the underlying muscular structure, i.e. due to the corresponding loss of structural integrity. Apart from the pure force–velocity relation, we also indicate the occurrence of the first significant skin injury in terms of IC2. This is indicated by the red triangle and was observed for each set of impacts for constant-mass configuration.

5.1. A.3. Three-dimensional plots

In order to generalize the obtained speed–mass relations and the injury occurrence results, it is convenient to evaluate three-dimensional surfaces in the speed, force, and mass space, obtained from the lines given in Figure 28. The surfaces depicted in Figure 29 represent every possible combination of velocity obtained for varying forces and masses under the assumption of linearity prior to muscle rupture, as suggested by the experimental data. The black spheres indicate the observed ‘key impacts’. An example of maximum speed–mass variation for a conservative constant force threshold is depicted by a red line. This function may e.g. serve as a threshold surface of ‘safe’ mass–velocity pairs. This, however, shows also that selecting a constant safe force independently of velocity and mass is a rather conservative threshold that does not capture the inherent structure of the results.

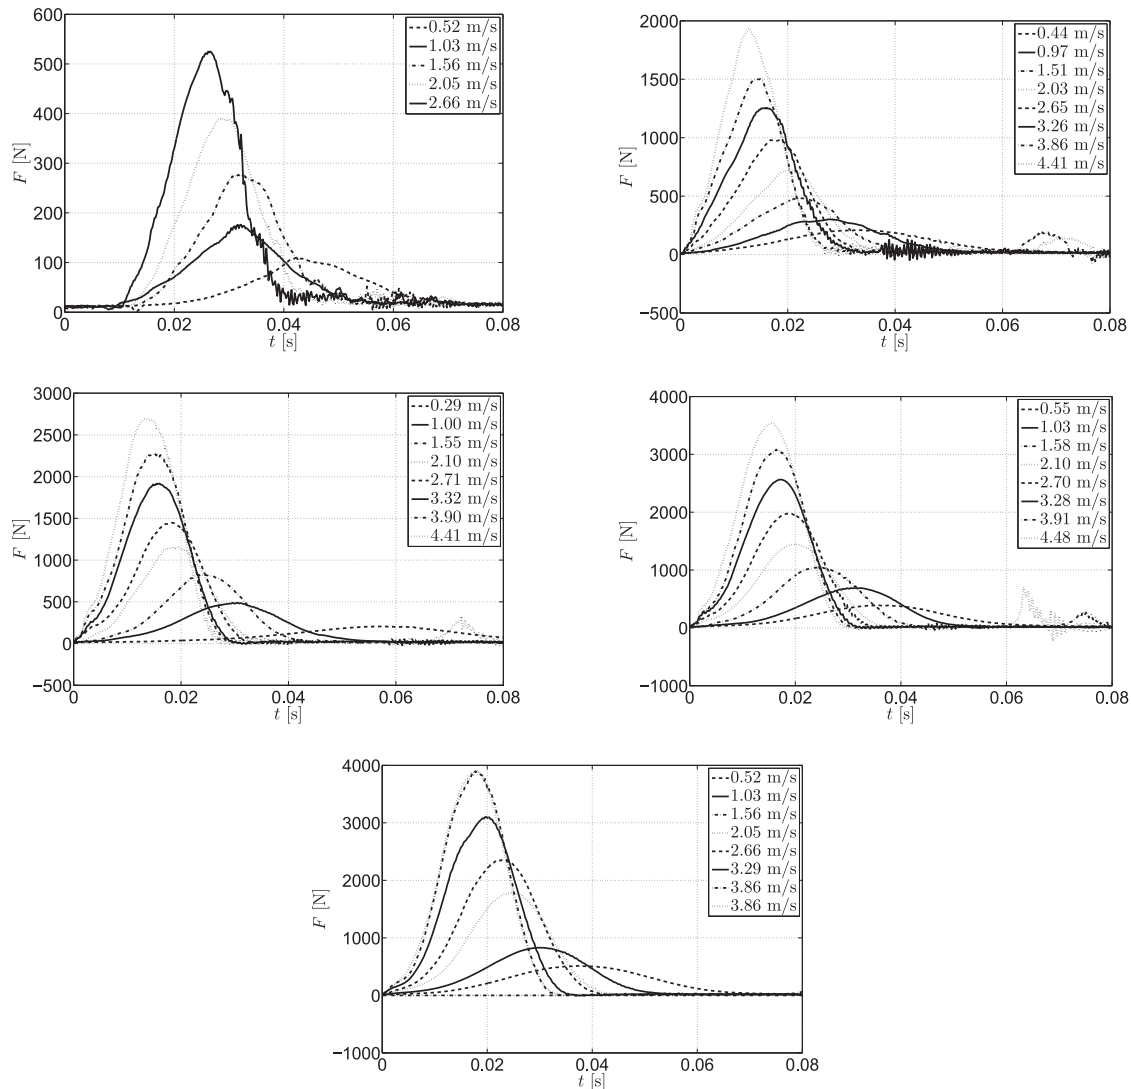


Fig. 25. Contact forces of the wedge for varying impact speed and mass (2.7 kg: upper left, 4.7 kg: upper right, 6.7 kg: middle left, 8.7 kg: middle right, 10.7 kg: bottom).

Next, we analyze pressure measurements, as it was recently proposed by BG/BGIA (2009) to be necessary for robot qualification.

A.4. Pressure evaluation

The behavior of biological soft-tissues during impacts is a complex process. For better understanding and prediction of the onset of injury it seems useful to consider the impact pressure as a relevant quantity. BG/BGIA (2009) recommended that pressure foils such as the Prescale Fuji Photo Film pressure foil are appropriate sensors to be applied. However, the results detailed next clearly indicate that for different reasons the use of these pressure foil is unsuitable for an analysis of general soft-tissue injury in most relevant cases. The main argument is its tendency to crumple especially during impacts of three-dimensional surfaces. As the

foil follows the shape changes imposed by the impactor on the soft-tissue during the collision, the measurement mainly consists of internal tension due to the inelasticity of the foil during the highly elastic soft-tissue collision,¹³ see Figure 30. This behavior is particularly evident in tests made with the $R = 12.5$ mm sphere. From Figure 30 one can see a star-shaped footprint that is of course not reasonable from a physical standpoint. However, as already mentioned it is important to underline that this problem depends on the dimension and the shape of the impactor. In particular, the pressure image of wedge impacts are reasonable. Furthermore, the footprint obtained with the smaller sphere ($R = 5$ mm) is rather consistent by means of shape. Unfortunately, saturation of the foil is another issue that makes these devices not very well applicable for our purposes of conducting large-scale experiments of biological tissue. As depicted in Figure 31 even though the measured forces

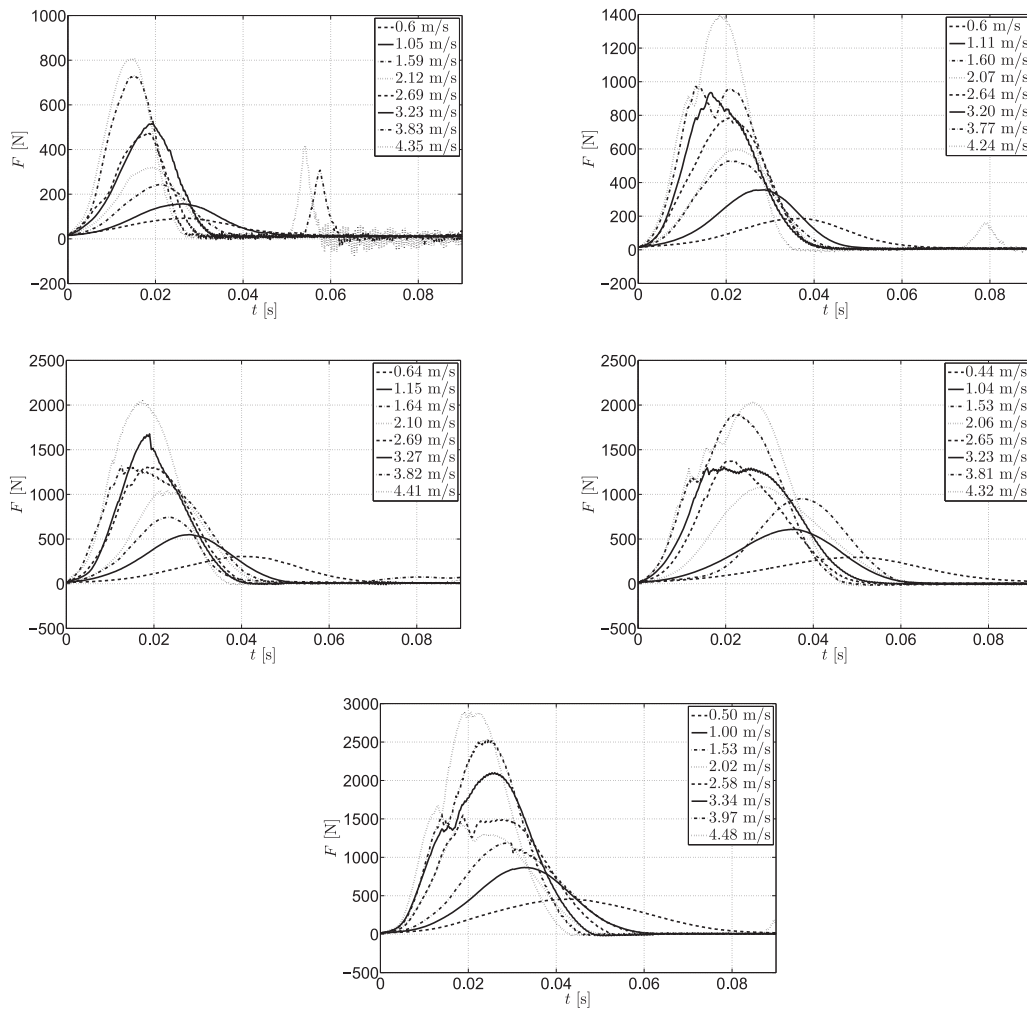


Fig. 26. Contact forces of the large sphere for varying impact speed and mass (2.2 kg: upper left, 4.2 kg: upper right, 6.2 kg: middle left, 8.2 kg: middle right, 10.2 kg: lower).

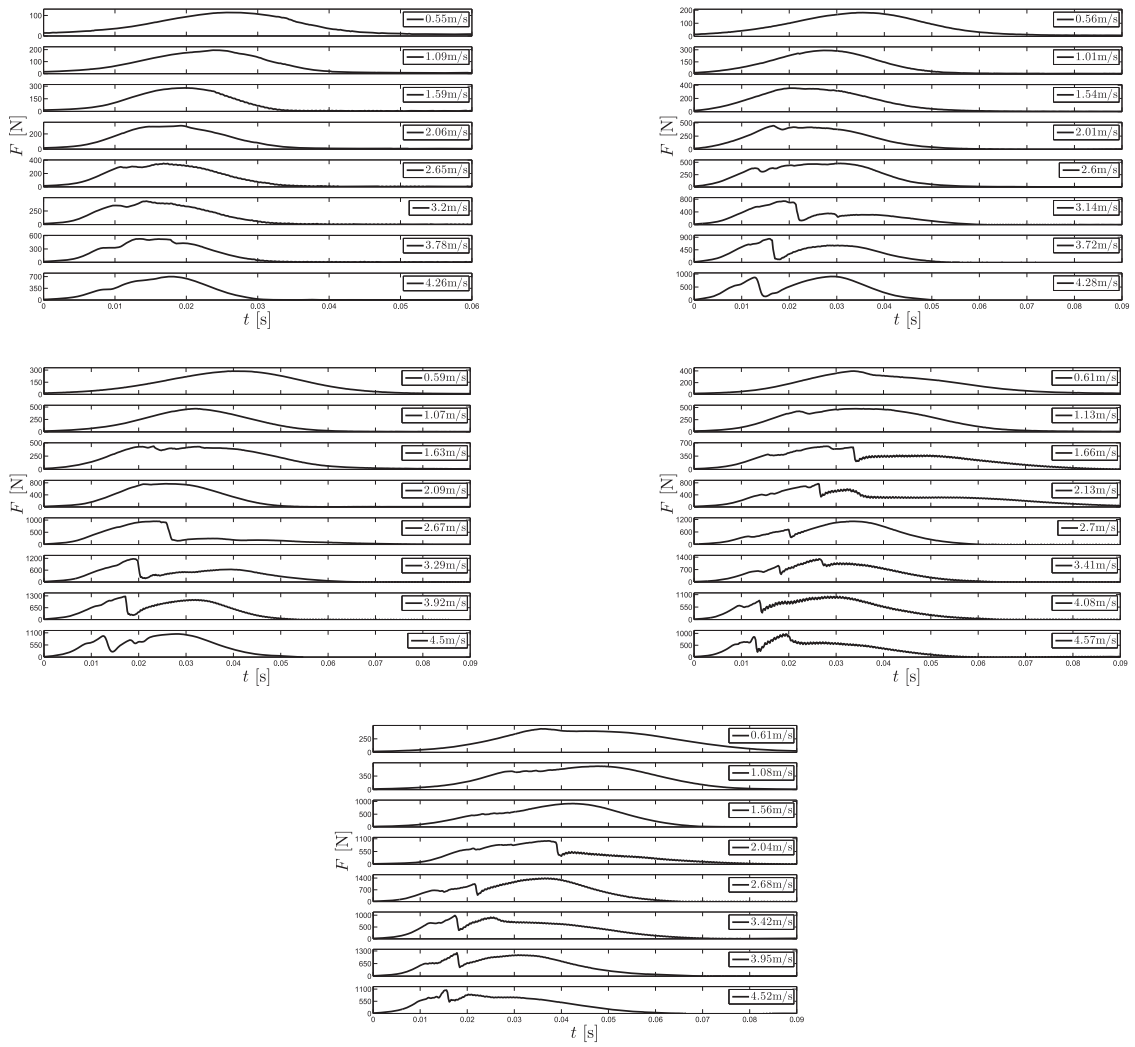


Fig. 27. Contact forces of the small sphere for varying impact speed and mass (2.1 kg: upper left, 4.1 kg: upper right, 6.1 kg: middle left, 8.1 kg: middle right, 10.1 kg: lower).

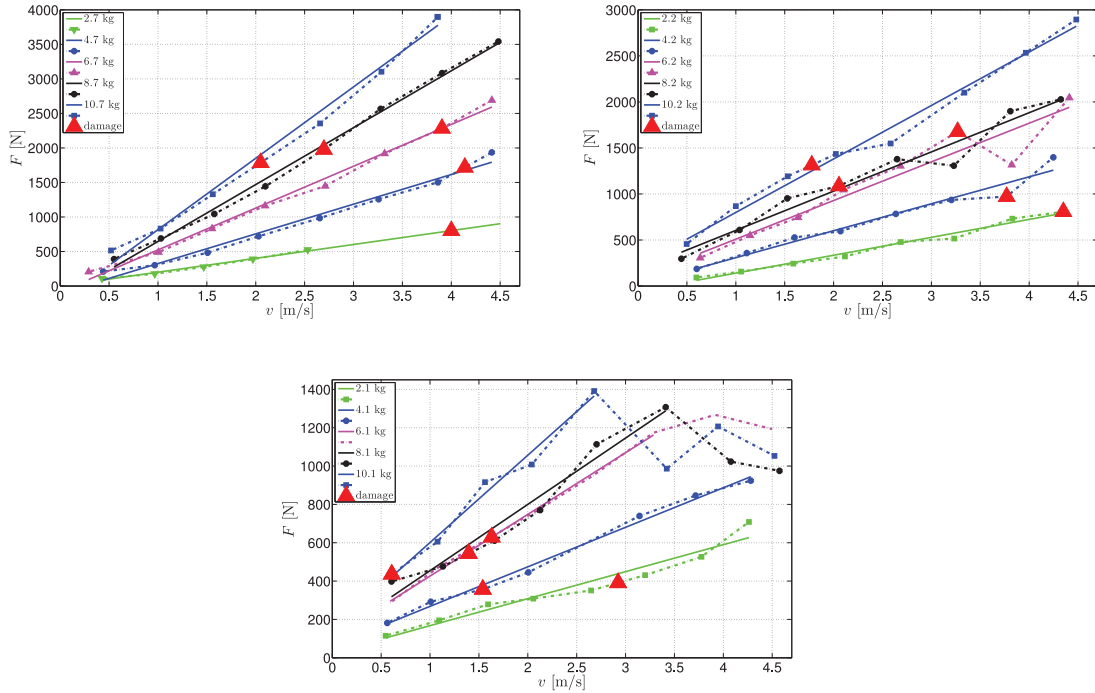


Fig. 28. Impact forces against impact velocity. The left plot shows the results for the wedge, while the right and bottom plots show the forces for the large and small sphere, respectively.

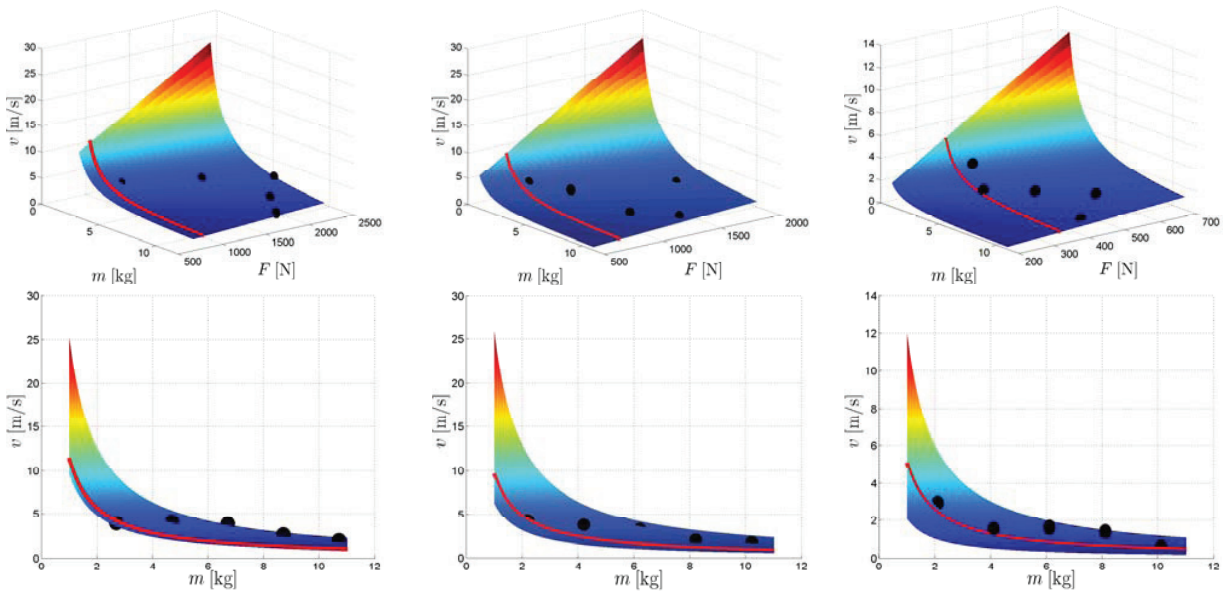


Fig. 29. Three-dimensional plot (upper row) and two-dimensional projection (lower row) of force, mass, and velocity. The black spheres are observed injuries and the red lines an example of speed–mass variation at constant force. The left column shows the results for wedge, while the middle and right one show the forces for the large sphere and the small sphere, respectively.

increase with velocity, the evaluated pressure remains almost constant for the small sphere, which is again physically not consistent.

To sum up, we can conclude that only the pressure evaluation of the wedge seems to be accurate and physically consistent with the measured contact forces.

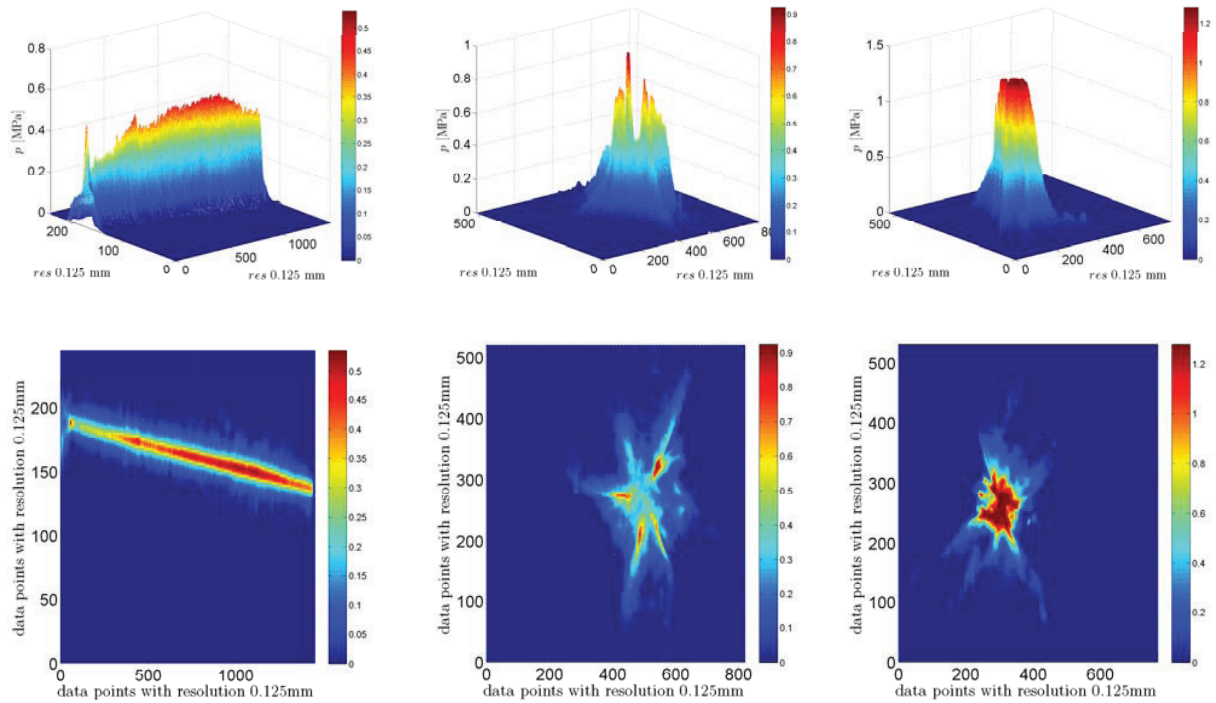


Fig. 30. Impact pressure at 1.5 m/s for the wedge 45° (left), large sphere (middle), and small sphere (right). The upper row shows the three-dimensional visualization and the lower row shows the impact profile from above. It is easy to observe that only the wedge (line impact) can be measured physically consistent and the pressure profiles of the spherical impactors (two-dimensional curvature) are characterized by internal tension caused by the incapability of the foil to bend in two dimensions.

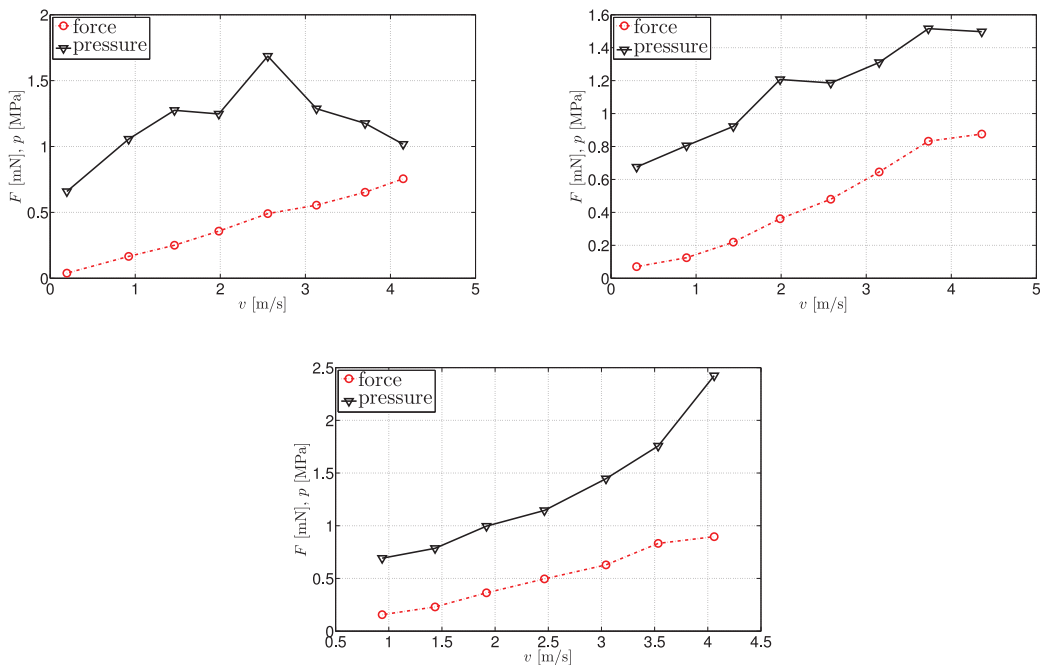


Fig. 31. Force/pressure over velocity for small sphere (left), large sphere (right), and wedge 45° (bottom). Again physical inconsistency can be observed due to the measurement principle of the sensor foil. This is especially the case for the small sphere.