



# Episodic Memories for Safety-Aware Robots

## Knowledge Representation and Reasoning for Robots that Safely Interact with Human Co-Workers

Georg Bartels<sup>1</sup> · Daniel Beßler<sup>1</sup> · Michael Beetz<sup>1</sup>

Received: 15 August 2018 / Accepted: 28 March 2019 / Published online: 9 April 2019  
© Gesellschaft für Informatik e.V. and Springer-Verlag GmbH Germany, part of Springer Nature 2019

### Abstract

In the factories and distribution centers of the future, humans and robots shall work together in close proximity and even physically interact. This shift to joint human–robot teams raises the question of how to ensure worker safety. In this manuscript, we present a novel episodic memory system for safety-aware robots. Using this system, the robots can answer questions about their actions at the level of safety concepts. We built this system as an extension of the KnowRob framework and its notion of episodic memories. We evaluated the system in a safe physical human–robot interaction (pHRI) experiment, in which a robot had to sort surgical instruments while also ensuring the safety of its human co-workers. Our experimental results show the efficacy of the system to act as a robot’s belief state for online reasoning, as well as its ability to support offline safety analysis through its fast and flexible query interface. To this end, we demonstrate the system’s ability to reconstruct its geometric environment, course of action, and motion parameters from descriptions of safety-relevant events. We also show-case the system’s capability to conduct statistical analysis.

**Keywords** Robot safety · Knowledge representation · Cognitive human–robot interaction

## 1 Introduction

In the factories and distribution centers of the future, humans and robots shall work together in close proximity and even physically interact. Such joint human–robot teams are an attractive proposition [13, 18]: companies hope to increase profits by automating stereotypical tasks with cheap robots. Furthermore, in periods of high work-load extra human co-workers can be flexibly added to boost throughput without incurring further investment costs. Finally, the humans’ jobs become less dangerous and more varied as the robots take over non-ergonomic and repetitive tasks.

From a technical point of view, this shift to joint human–robot teams introduces severe challenges. Most pressing, there is the question of worker safety [2, 18]: how

to protect them from injury and worse despite the physical strength of the robots?


In our own work, we focus on developing the cognitive capabilities of what we call *safety-aware robots* [6]. That is, robots that know when their actions have the potential to hurt or threaten a human co-worker, and that actively refrain from performing such actions [6]. A key cognitive capability of safety-aware robots is the ability to remember actions and conceptualize them in terms of safety-relevant events.

Using such a memory system, the robots can answer questions about their actions at the right level of abstraction, i.e. the level of safety concepts. Example queries are “Was I close to a vulnerable body part of my co-worker?” or “What type of collision events occurred during pick-up actions?”<sup>1</sup> This type of question answering capability is a crucial resource for both run-time decision making and offline safety analysis.

✉ Georg Bartels  
georg.bartels@uni-bremen.de

<sup>1</sup> Institute for Artificial Intelligence (IAI), University of Bremen, 28359 Bremen, Germany

<sup>1</sup> The memory system that we present in this system does not process natural language questions. Instead, it processes logic queries. We use natural language questions as an illustration of the expressivity of the query language.



**Fig. 1** Experimental setup of the evaluation scenario: a table-mounted robot shall sort surgical tools into a basket while ensuring the safety of human co-workers that can physically interact with the robot and the tools, at any moment

In this manuscript, we present a novel episodic memory system for safety-aware robots. The system encodes the actions and percepts of safety-aware robots in detail, and provides a fast and flexible query interface for online reasoning and offline safety analysis. We built the system as an extension of the *KNOWROB* framework [23] and its notion of episodic memories [7].

To evaluate the episodic memory system, we integrated it into a robot that sorted surgical instruments for medical operations. Throughout its work, the robot had to ensure the safety of humans that could physically interact with the tools and the robot. Figure 1 depicts the experimental setup. During the experiments, the robot kept track of all surgical instruments and recovered from human interruptions using online reasoning on its episodic memories. We performed 15 experimental runs with an overall length of 47min. During those experiments, the robot performed 197 pick-and-place actions while safely handling 196 intrusion events and 226 contact events; show-casing the efficacy of the overall system. After the experiments, we used the system to conduct a safety analysis of the robot's behavior. The results show that the system is able to reconstruct its geometric environment, course of action, and motion parameters from descriptions of safety-relevant events.

The contributions of this manuscripts are: (1) we present an extension of the *KNOWROB* framework and its notions of episodic memories for safety concepts in physical human–robot interaction (pHRI). (2) We show that the episodic memory system can be used for online reasoning, using a pHRI experiment in which a robot sorts surgical instruments. (3) Using the episodic memories from the same experiments, we show-case the system's capability to reconstruct the robot's geometric environment, course of action,

and motion parameters from descriptions of safety-relevant events.

The remainder of this manuscript is structured as follows: The Sect. 2 reviews the related work, and Sect. 3 describes our software architecture and presents our safety-related extension to *KNOWROB* and its episodic memory system. Section 4 evaluates the episodic memory system, and Sect. 5 concludes the manuscript.

## 2 Related Work

In most production lines, robots are confined to safety cages which represent an effective and efficient solution to the question of worker safety. However, by definition this approach is incompatible with the idea of human and robotic co-workers sharing a work-space [13, 18]. In recent years, several companies have introduced robots that shall forego their safety cages<sup>2,3,4</sup>. The safety concepts of this new generation of robots typically comprise a variety of mechanical and motor control safety features, such as lightweight constructions with low-power motors [14], impedance control for soft contact behavior [1], or external disturbance observers that trigger controlled stops upon collisions [12, 19].

These developments have endowed the robotics community with a rich and safe motion generation apparatus for robots. Now, the research question is how to properly configure these modules at a higher level of abstraction [6]? Typically, these modules make fast decisions using instantaneous sensors readings such as joint angle or torque measurements that, unfortunately, often do not reveal the underlying semantics of an interaction, e.g. “Has the robot collided with the head or thigh of the human?” Here, emerging technologies from human detection, localization, and tracking have opened up exciting new research opportunities [11].

Recently, several systems for planning safe motions for human–robot interactions have been presented [21, 26]. These systems typically employ depth sensors to track humans in the vicinity of the robot to plan safe robot motions [21, 26]. While these studies focused on avoiding pHRI, we investigate how to represent and reason about safe pHRI episodes.

A recent survey on safe HRI divides the field into four sub-fields: Safety through control, motion planning, prediction, and psychological consideration [16]. Interestingly, our work does not really fit into either category. Here, we present a knowledge base for safety-aware robots that encodes the

<sup>2</sup> <https://www.rethinkrobotics.com/sawyer>.

<sup>3</sup> <https://www.kuka.com/en-de/products/robot-systems/industrial-robots/lbr-iiwa>.

<sup>4</sup> <https://www.franka.de/panda>.

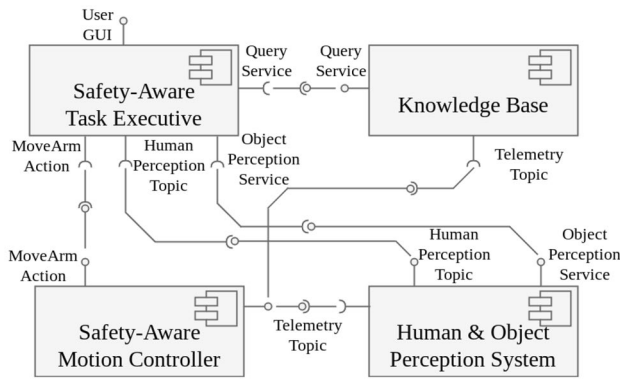


Fig. 2 UML diagram of our ROS software architecture

activities in pHRI scenarios using episodic memories. This knowledge base allows the robot and its human developers to analyze the interactions in terms safety events such as intrusions, collisions, etc. The only other robot knowledge base for HRI that we are aware of is presented in [17]. However, that systems focused on dialog grounding, and neither considered physical interactions nor safety aspects.

### 3 Methodology

#### 3.1 Software Architecture

Figure 2 depicts the software architecture of our robotic system as an UML diagram. It consists of four components: a safety-aware task executive, a knowledge base, a safety-aware motion controller, and a human and object perception system. We developed and integrated the components using the Robot Operating System (ROS) middleware [20].

The safety-aware task executive is the main entry point into the application. It allows users to start and stop the robot through a GUI, and coordinates the behavior of the other software components. This executive has been implemented using the cognitive robot abstract machine (CRAM) framework<sup>5</sup> [3], following the paradigm of cognition-enabled robot control [4, 9].

The safety-aware motion controller combines a suite of control schemes that move the robot, e.g. Cartesian and joint Impedance [1], with a model-based disturbance observer to detect collisions [12]. Using this combination, it is possible to implement fast safety reactions that quickly stop the robot in a controlled manner upon detection of collisions [12]. In this experiment, we used the system described by Parusel et. al. [19].

<sup>5</sup> <http://cram-system.org>.

The perception system uses a wrist-mounted RGB sensor to detect, recognize, and localize the surgical instruments on the table, and employs a wall-mounted RGB-D sensor to detect, and localize the humans around the robot<sup>6</sup>. It is implemented using the RoboSherlock framework<sup>7</sup> [5].

The knowledge base encodes the system’s belief state as episodic memories. It gives other components access to the belief state through its query interface. The knowledge base was realized as an extension of the KnowRob framework<sup>8</sup> [23] that employs the W3C Web Ontology Language (OWL) [10] for knowledge representation and SWI-Prolog [24] as its query interface. The details of this component are described in the remainder of this article.

For more information on our software architecture, in particular the safety-aware task executive and the perception system, we refer the reader to one of our prior publications that describes an earlier version of the system [6]. In the remainder of this manuscript, we will focus our description on the knowledge base as it constitutes the main contribution of this publication.

#### 3.2 A Knowledge Base for Safety-Aware Robots

The knowledge base for our safety-aware robots extends the KnowRob framework and its notion of episodic memories [7, 8]. Episodic memories are time-indexed recordings of sub-symbolic data coupled with a symbolic narration of the activity<sup>9</sup>. Here, we concentrated on extending the existing representations of episodic memories of KnowRob with concepts for safe pHRI scenarios. That involves concepts such as monitoring humans, human intrusions, and physical contact with humans. In particular, we incorporated a notion of human co-workers, the poses of their body parts, and safety-relevant events. The rest of this section will describe these in more detail.

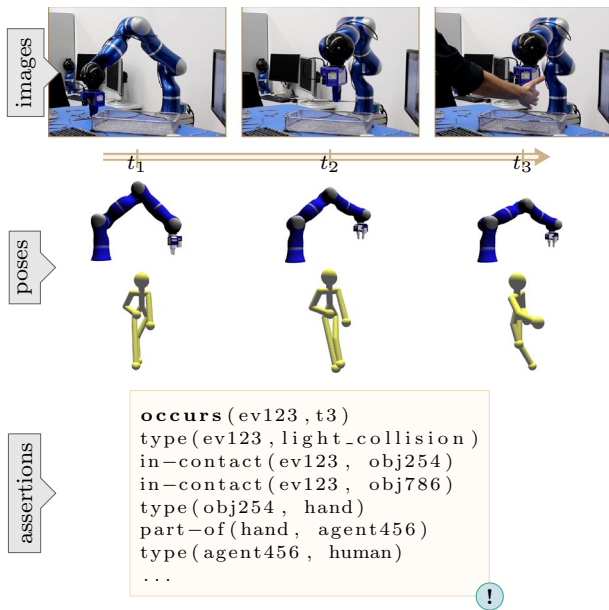
Figure 3 depicts an exemplary episodic memory with safety features. It shows the narration of a collision event  $ev_{123}$  that occurred at time instant  $t_3$  between the hand of a human and some other object, and also the pose of the human and the robot at that time instant. Using the episodic memory system, it is possible to retrieve and reconstruct the episode and visualize it for offline analysis. Figure 4 depicts

<sup>6</sup> The wall-mounted RGB-D sensor is not a safety-rate sensor. It was used to demonstrate the use of the knowledge base and not to demonstrate the efficacy of the RGB-D sensor as a safety-rated sensor.


<sup>7</sup> <http://robosherlock.org>.

<sup>8</sup> <http://knowrob.org>.

<sup>9</sup> In our experiments, we used the CRAM extension described in [25] to automatically encoded the episodic memories from the safety-aware task executive.



**Fig. 3** Illustration of episodic memories of safety aware robots. Top: Sub-symbolic data in the form of images captured and poses recorded over time. Bottom: Symbolic assertions that describe the activity at a high level of abstraction



**Fig. 4** Visualization of the reconstructed interaction depicted in Fig. 3 using the episodic memory system

such a reconstruction for the episode which is displayed in Fig. 3<sup>10</sup>.

### 3.2.1 Human Tracking and Representation

The basis of safe pHRI is that robots are aware of the humans in their vicinity. Our episodic memory system encodes two types of information about humans. Firstly, it encodes that a particular human individual is present and which of its body parts were perceived at a given moment in time. This information is encoded symbolically w.r.t. the system's ontology. Secondly, the system continuously encodes the

spatial location of all perceived body parts as sub-symbolic data using homogeneous transforms, i.e. 3D translation and orientation.

The symbolic representations encode which body parts of the humans are tracked at a given moment in time and what type these parts have, e.g. left shoulder, right forearm, or head. Using these asserted links and joints, we derive semantic components that are defined as kinematic chains of links and joints that have associated semantics. For instance, an arm is a component from shoulder to hand base, and arm components are used for manipulation activities. To this end, we adopted the semantic robot description language (SRDL) of KnowRob [15] to also describe the kinematic structure and capabilities of humans.

Regarding the sub-symbolic human representations, each body part individual is assigned a coordinate frame with a name that is identical to its own name<sup>11</sup>. Using this assumption and a shared timestamp index, the system can retrieve the pose of each human body part at every moment in time. As a result, we can use symbolic activity descriptions as an index to retrieve human pose data. Additionally, it is also possible to go the other way, i.e. from sub-symbolic to symbolic representations, and to compute event symbols such as situations when a human got too close to the robot, on demand.

During our experiments, the number of humans is not pre-determined, and new unique symbols describing the perceived humans and their body parts must be grounded ad-hoc in the data structures of the human tracking system. This is supported by the perception system that tracks the identity of humans over time.

### 3.2.2 Safety-Relevant Event Representation

There are three different types of safety-relevant events encoded in the episodic memories: Human perception events, human intrusion events, and contact events.

Whenever a human is detected, an event is generated by the tracker that is also represented symbolically in the episodic memories. KnowRob comes with a notion of *Perception* which is an event caused by the perception system in which one or more objects were *detected*. The relation *detected* is defined as *post-actor* in terms of KnowRob's action ontology [22] (i.e., post-condition). We define a new subclass of *Perception*, called *PersonPerception*, and further specify that only persons are detected in *PersonPerception* events. The complex description of detected persons

<sup>10</sup> Visual scene reconstruction is just one of the possible applications of our system, albeit a very intuitive one.

<sup>11</sup> The symbolic human description we use is tailored to the modeling that the tracker provides. For instance, the tracker only detects one link when monitoring the hand of a human. Hence, human hands are modelled with only one link.

(including semantic components, links, and joints) is then generated using a template SRDL ontology, where the id of the human is prefixed to the names of OWL individuals and coordinate frames.

Whenever a human gets too close to the robot, humans may get injured if the robot continues its operation. We say that entering the close proximity of the robot is an intrusion into its direct workspace. The detection of intrusions is performed by the plan executive based on the distance of monitored humans to the tool frame of the robot. Intrusions are detected when this distance is below a certain threshold (0.4m for a person's head, and 0.3m for her hands). In the terminological knowledge base, we define a concept *HumanIntrusion* to represent these events. Body parts causing intrusion events are linked to the corresponding event via *pre-actor* assertions (i.e., pre-conditions).

During the experiment a sub-routine of the plan executive is also monitoring collision events. We define thresholds for the classification of four different collision types: "severe collision", "strong collision", "light collision", and "contact"<sup>12, 13</sup>. Finally, symbols are asserted that describe the detected event types.

### 3.2.3 Task Representation and Online Reasoning

Our system generates its episodic memories during run-time on the fly, and can also access them for online reasoning. The safety-aware plan executive extensively uses this feature. It encodes each of its actions and their resulting effects in the belief state. Afterwards, it uses the same belief state to infer suitable plan parameters based on past actions and experiences.

For example, the belief state contains descriptions of detected objects and persons, and how they relate to actions performed by the robot and events that were detected. In our scenario, the task of the robot is to put a specific set of items into a basket at pre-defined locations. We describe this as a kind of "recipe" in the terminology of our knowledge base by decomposing the basket into different parts which are slots with a transform relative to the basket, and which are the storage place for a specific object type. The state of a slot is inferred from the actions that were performed by the robot. A slot is inferred as occupied if a *putting down* action with a target location that is identical to the slot's location was successfully executed.

Given this information, the robot compares the detected objects with the current state of the basket to infer candidate objects for empty slots in the baskets, or where to put an object it currently holds. It can also conclude that the *find-object* sub-routine must be activated in case there are empty slots left but no suitable perceived object is asserted in the episodic memory.

## 4 Evaluation

### 4.1 Experimental Setup

We evaluated the episodic memory system by deploying it in a safe pHRI experiment. As robot we employed a table-mounted DLR LWR3 robot [14] equipped with Weiss WSG 50 industrial gripper.<sup>14</sup> The robot was also equipped with a wrist-mounted RGB camera to perceive surgical instruments on the table. Additionally, there was a wall-mounted RGB-D sensor to perceive humans in the vicinity of the table.

The robot had to sort the surgical instruments on the table into a basket that was also located on the table. To this end, the robot had to perceive the instruments on the table, pick them up, and place them into the basket. All together, there were 8 types of instruments, such as scalpels, scissors, forceps, and pincers. The multiplicity of the instruments on the table were not known beforehand to the robot, i.e. some objects could be missing or be present multiple times.


During sorting, the robot had to ensure the safety of its human co-workers, and recover from their disturbances. Whenever a human got too close to the robot, the system recorded an intrusion event, stopped, and waited for the human to retreat again. Human participants were allowed to rearrange the objects on the table, and even to remove or add some of them. Additionally, participants were allowed to physically interact with the robot and take instruments from its gripper. During pHRI, the robot would log a contact event and trigger a specified safety reaction, e.g. controlled stop, stop with brakes, or switching into gravity compensation mode. After the end of intrusion and collision events, the task executive would issue the required recovery actions, such as re-grasping or re-perceiving of instruments, and resume sorting.<sup>15</sup>

<sup>12</sup> The safety-aware motion controller reports these collision types whenever the measured external torque of any joint exceeds 5%, 10%, 15%, and 30% of the joint's maximum torque, respectively.

<sup>13</sup> For this experiment, we chose threshold values and collision classes on the basis of informed guesses. In the future, we plan to comply with recent industry standards, e.g. ISO/TS 15066 (<https://www.iso.org/standard/62996.html>).

<sup>14</sup> <https://weiss-robotics.com/en/produkte/gripping-systems/performance-line-en/wsg-50-en>.

<sup>15</sup> Obviously, this is not a collaborative human–robot interaction experiment. At best, human co-workers would not interfere with the robot. In most cases, the interaction would disturb the robot and it had to safely cope.



**Fig. 5** Different reconstructed scenes: Co-occurring contact and intrusion events (top left), a contact event without co-occurring intrusion events (top right), a trajectory of a human during an intrusion event (bottom left), and all instruments perceived by a perception action (bottom right)

Figure 1 displays our experimental setup with the robot, surgical tools, and a human reaching for the robot, causing an intrusion event.

## 4.2 System Efficacy and Online Reasoning

All together, we performed 15 experimental runs with an overall length of 47 min. Within those 15 episodic memories, the robot encoded 197 pick-and-place actions, 40 human tracking actions, 196 intrusion events, and 226 contact events. We conducted the experiments during the final review of the SAPHARI project,<sup>16</sup> with roughly 3 dozen external experts as audience. This experimental evaluation showed that the integrated system was effective in sorting surgical instruments while at the same time ensuring the safety of its human co-workers. Furthermore, the experiments show that the system was effective in acting as the robot’s belief state and online reasoning resource. Video recordings of the experiments, as well as the episodic memories themselves are available through our web-based knowledge service OPENEASE<sup>17</sup>

## 4.3 Event Retrieval

To evaluate the episodic memory system itself, we employ its query interface to retrieve and visualize various events from the experiments. This evaluation show-cases the system’s capability to reconstruct the robots geometric

environment, course of action, and motion parameters using descriptions of safety-relevant events. Figure 5 depicts reconstructed the scenes from four evaluation queries that we discuss in this sub-section.

Using Prolog, users can conveniently search the episodic memories using descriptions of co-occurring events. For instance, to have the system visualize the scene when contact and intrusion events occurred at the same time, one can formulate this query:

```
?- entity(Intrusion, [an,event,[type,intrusion]]),
   entity(Contact, [an,event,[type,contact]]),
   holds(during(Contact,Intrusion)),
   show(Contact).
```

Furthermore, it is also possible to express that certain events should not occur during moments of interest. As an example, consider the following query that displays the scene for contact events that happened without co-occurring intrusion events:<sup>18</sup>

```
?- entity(Contact, [an,event,[type,contact]]),
   \+ (
     entity(Intr, [an,event,[type,intrusion]]),
     holds(during(Contact,Intr))),
   show(Contact).
```

It is also possible to retrieve and display trajectories of body parts for both the robot and its human co-workers. The following query displays the scene at the start of an intrusion event and plots the trajectory of the intruding human body part using red spheres:

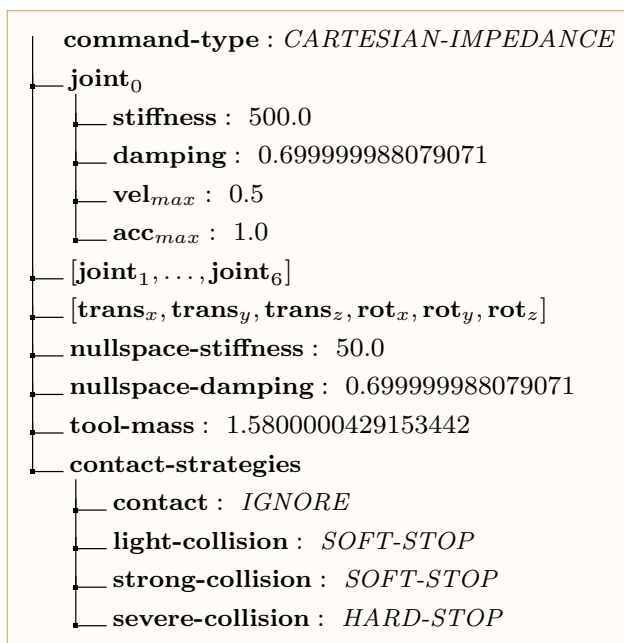
```
?- entity(Intrusion, [an,event,[type,intrusion],
                    [body-part-used,Part]]),
   occurs(Intrusion, [Begin,End]),
   show(trajectory(Part), [Begin,End]).
```

Besides encoding events like contact and intrusion events, the system also records its actions and their resulting events. Let’s consider perception events as an example. The next query retrieves the perceptual results of an instrument perception action and displays the perceived scene in the canvas:

<sup>16</sup> <http://www.saphari.eu>.

<sup>17</sup> <http://www.open-ease.org> please select experiment “Safe Interaction” and then “Preparing instruments”.

<sup>18</sup> Such events typically occurred during pick-up actions when contact between gripper and table caused the safety-aware motion controller to raise a contact event.



**Fig. 6** An action designator retrieved from the episodic memory that describes a desired motion as commanded to the safety-aware motion controller

```


?- entity(Act, [an, action,
               [type, voluntary_body_movement],
               [designator, D]]),
   holds(knowrob:equatedDesignator(D,E)),
   show(E).
    
```

Figure 6 depicts one such action designator retrieved from the episodic memories. The retrieved designator gives detailed insight into the commanded motion parameters. It specifies the control scheme, stiffness and damping parameters, velocity and acceleration thresholds, tool configuration parameters, as well as the specified safety reactions in case of collisions.<sup>19</sup>

### 4.4 Episode Statistics

In addition to retrieving individual events, the episodic memory system can also be used to perform statistical analysis of entire episodes. To this end, users can formulate queries to find all events that match particular descriptions and can employ built-in visualization charts, such as pie charts, time lines, or histograms.

The following two queries demonstrate this capability using pie chart diagrams. The resulting query visualization are depicted in Fig. 7. The first query (depicted on the left side of Fig. 7) retrieves all contact events that occurred in the experiments, and displays the distribution of contact event types as a pie chart:



**Fig. 7** The distribution of contact event types (left) and durations of human intrusions (right) over 15 runs. 196 intrusions and 226 contacts occurred in total

```

?- findall(EvT, (
   entity(Ev, [an, event, [type, contact]]),
   type_of(Ev,EvT)), EventType),
   data_vis(piechart, [data:EventType]).
    
```

The right side of the Fig. 7 depicts the result of a more complex query. Here, all human intrusion events are retrieved and their respective durations are grouped and counted into three user-defined ranges. Namely, intrusions with at most 3s duration, intrusions with duration 3s to 6s, and intrusions with at least 6s duration. The results are again depicted as a pie chart:

```

?- entity(Perc, [an, action, [type, perception],
               [perception_result, Obj]],
   occurs(Perc, [Instant, -]),
   show(Obj, Instant).
    
```

During sorting, the task executive also encodes the motion commands it sends to the safety-aware motion controller. These commands are not stored as axiomatic knowledge, but rather as *action designators*, i.e. flexible lists of key value pairs in the sub-symbolic part of the episodic memories. To enable retrieval, those designators are directly linked to the encoded symbolic actions. As an example, the following query retrieves the designators of all actions of type *voluntary\_body\_movement*, a rather abstract superclass of many actions such as grasping, reaching and placing:

<sup>19</sup> The designator contains various constants, e.g. *CARTESIAN-IMPEDANCE*, *light-collision*, *HARD-STOP*. These are pre-defined constants of the communication protocol between the safety-aware task executive and the safety-aware motion controller.

```
?- Ranges = [[0.0,3.0,'0-3s'],
             [3.0,6.0,'3-6s'],
             [6.0,9999.0,'>6s']],
   findall(Label, (
     entity(Ev, [an, event, [type, intrusion]]),
     occurs(Ev, [Begin,End]),
     Duration is End - Begin,
     member([Min,Max,Label], Ranges),
     Min <= Duration,
     Max >= Duration),
   RangeLabels),
   data_vis(piechart, [data:RangeLabels])).
```

## 5 Conclusion

In this manuscript, we presented a novel episodic memory system for safety-aware robots. The system was built as an extension of the KnowRob framework and its notion of episodic memories. We evaluated the system in a safe pHRI experiment in which a robot had to sort surgical instruments while ensuring the safety of its human co-workers. Our experiments showed the efficacy of the episodic memory system to act as the robot's belief state system for online reasoning. Furthermore, we show-cased the system's capability to reconstruct its geometric environment, course of action, and action parameterization from descriptions of safety-relevant events. Finally, we also showed how the episodic memory system can be employed to perform statistical analysis of entire pHRI episodes.

**Acknowledgements** We gratefully acknowledge that this work was partially funded by the FP7 Project *SAPHARI* (Project ID: 287513) and by Deutsche Forschungsgemeinschaft (DFG) through the Collaborative Research Center 1320 *EASE*.

## References

1. Albu-Schäffer A, Ott C, Hirzinger G (2007) A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *Int J Robot Res* 26:23–39
2. Badri A, Boudreau-Trudel B, Souissi AS (2018) Occupational health and safety in the industry 4.0 era: A cause for major concern? *Safety Sci* 109:403–411
3. Beetz M, Mösenlechner L, Tenorth M (2010) Cram—a cognitive robot abstract machine for everyday manipulation in human environments. In: *International conference on intelligent robots and systems IEEE*, pp 1012–1017
4. Beetz M, Jain D, Mosenlechner L, Tenorth M, Kunze L, Blodow N, Pangercic D (2012) Cognition-enabled autonomous robot control for the realization of home chore task intelligence. *Proc IEEE* 100:2454–2471
5. Beetz M, Bálint-Benczédi F, Blodow N, Nyga D, Wiedemeyer T, Marton ZC (2015a) Robosherlock: Unstructured information processing for robot perception. In: *International conference on robotics and automation IEEE*, pp 1549–1556
6. Beetz M, Bartels G, Albu-Schäffer A, Bálint-Benczédi F, Belder R, Beßler D, Haddadin S, Maldonado A, Mansfeld N, Wiedemeyer T, et al (2015b) Robotic agents capable of natural and safe physical interaction with human co-workers. In: *International conference on intelligent robots and systems IEEE*, pp 6528–6535
7. Beetz M, Tenorth M, Winkler J (2015c) Open-ease. In: *International conference on robotics and automation IEEE*, pp 1983–1990
8. Beetz M, Beßler D, Haidu A, Pomarlan M, Bozcuoglu AK, Bartels G (2018) Knowrob 2.0 – a 2nd generation knowledge processing framework for cognition-enabled robotic agents. In: *International conference on robotics and automation IEEE*, pp 512–519
9. Ersen M, Oztop E, Sariel S (2017) Cognition-enabled robot manipulation in human environments: requirements, recent work, and open problems. *Robot Autom Mag* 24:108–122
10. OWL Working Group et al (2009) OWL 2 web ontology language document overview: W3C recommendation 27 October 2009. <https://www.w3.org/TR/2009/REC-owl2-overview-20091027>
11. Haddadin S, Croft E (2016) *Physical human–robot interaction*. In: *Springer handbook of robotics*. Springer, New York, pp 1835–1874
12. Haddadin S, Albu-Schäffer A, De Luca A, Hirzinger G (2008) Collision detection and reaction: A contribution to safe physical human–robot interaction. In: *International conference on intelligent robots and systems IEEE*, pp 3356–3363
13. Hägele M, Schaaf W, Helms E (2002) Robot assistants at manual workplaces: effective co-operation and safety aspects. In: *International Symposium on Robotics*, vol 7
14. Hirzinger G, Sporer N, Albu-Schäffer A, Hahnle M, Krenn R, Pascucci A, Schedl M (2002) Dlr's torque-controlled light weight robot iii—are we reaching the technological limits now? In: *International conference on robotics and automation IEEE*, vol 2, pp 1710–1716
15. Kunze L, Roehm T, Beetz M (2011) Towards semantic robot description languages. In: *International conference on robotics and automation IEEE*, pp 5589–5595
16. Lasota PA, Fong T, Shah JA et al (2017) A survey of methods for safe human–robot interaction. *Found Trends Robot* 5:261–349
17. Lemaignan S, Ros R, Mösenlechner L, Alami R, Beetz M (2010) Oro, a knowledge management platform for cognitive architectures in robotics. In: *International conference on intelligent robots and systems, IEEE*, pp 3548–3553
18. Matthias B, Kock S, Jerregard H, Källman M, Lundberg I (2011) Safety of collaborative industrial robots: Certification possibilities for a collaborative assembly robot concept. In: *International symposium on assembly and manufacturing IEEE*, pp 1–6
19. Parusel S, Haddadin S, Albu-Schäffer A (2011) Modular state-based behavior control for safe human–robot interaction: A lightweight control architecture for a lightweight robot. In: *International conference on robotics and automation IEEE*, pp 4298–4305
20. Quigley M, Conley K, Gerkey B, Faust J, Foote T, Leibs J, Wheeler R, Ng AY (2009) Ros: an open-source robot operating system. In: *ICRA workshop on open source software*, vol 3, pp 5–10
21. Ragaglia M, Zanchettin AM, Rocco P (2018) Trajectory generation algorithm for safe human–robot collaboration based on multiple depth sensor measurements. *Mechatronics* 55:267–281
22. Tenorth M, Beetz M (2012) A unified representation for reasoning about robot actions, processes, and their effects on objects. In: *International conference on intelligent robots and systems IEEE*, pp 1351–1358
23. Tenorth M, Beetz M (2013) Knowrob: a knowledge processing infrastructure for cognition-enabled robots. *Int J Robot Res* 32:566–590
24. Wielemaker J, Schrijvers T, Triska M, Lager T (2012) Swi-prolog. *Theory Pract Logic Program* 12:67–96
25. Winkler J, Tenorth M, Bozcuoglu A, Beetz M (2014) Crammemories for robots performing everyday manipulation activities. *Adv Cogn Syst* 3:47–66
26. Zanchettin AM, Ceriani NM, Rocco P, Ding H, Matthias B (2016) Safety in human–robot collaborative manufacturing environments: Metrics and control. *Trans Autom Sci Eng* 13:882–893