Computers in Industry 138 (2022) 103627

Contents lists available at ScienceDirect

Computers in Industry

journal homepage: www.elsevier.com/locate/compind

OCRA – An ontology for collaborative robotics and adaptation

Alberto Olivares-Alarcos[a,*], Sergi Foix[a], Stefano Borgo[b], Guillem Alenyà,[a]

[a] Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Consejo Superior de Investigaciones Científicas (CSIC) - Universitat Politècnica de Catalunya (UPC), Llorens i Artigas 4-6, 08028 Barcelona, Spain
[b] Laboratory for Applied Ontology (LOA), ISTC-CNR, Trento, Italy

ARTICLE INFO

Article history:
Received 3 September 2021
Received in revised form 28 December 2021
Accepted 27 January 2022
Available online xxxx

Keywords:
Ontology
Ontological analysis
Industrial collaborative robots
Human-robot collaboration
Adaptation
Reliability

ABSTRACT

Industrial collaborative robots will be used in unstructured scenarios and a large variety of tasks in the near future. These robots shall collaborate with humans, who will add uncertainty and safety constraints to the execution of industrial robotic tasks. Hence, trustworthy collaborative robots must be able to reason about their collaboration's requirements (e.g., safety), as well as the adaptation of their plans due to unexpected situations. A common approach to reasoning is to represent the knowledge of interest using logic-based formalisms, such as ontologies. However, there is not an established ontology defining notions such as collaboration or adaptation yet. In this article, we propose an Ontology for Collaborative Robotics and Adaptation (OCRA), which is built around two main notions: collaboration, and plan adaptation. OCRA ensures a reliable human-robot collaboration, since robots can formalize, and reason about their plan adaptations and collaborations in unstructured collaborative robotic scenarios. Furthermore, our ontology enhances the reusability of the domain's terminology, allowing robots to represent their knowledge about different collaborative and adaptive situations. We validate our formal model, first, by demonstrating that a robot may answer a set of competency questions using OCRA. Second, by studying the formalization's performance in limit cases that include instances with incongruent and incomplete axioms. For both validations, the example use case consists in a human and a robot collaborating on the filling of a tray.

1. Introduction

During the last decade, the industrial sector has shown a growing interest in more flexible manufacturing processes where humans and robots are expected to work together. For that purpose, collaborative robots, or co-bots, are robots specifically designed for direct interaction with humans within a collaborative workspace (ISO, 2011a). Implementing industrial processes where robots and humans collaborate, opens several questions such as how to cope with uncertainty and safety. Hence, collaborative robots shall be able to, among others, reason about their tasks' requirements (e.g. safety, performance, etc.), about the changes in their environment, and about the plan adaptations due to those changes.

The use of industrial collaborative robots has drawn the attention of many researchers, becoming a prolific research domain (Gervasi et al., 2020; Gualtieri et al., 2021; Kim et al., 2021). Indeed, several works have discussed safety in collaborative robotic scenarios (Vicentini, 2020; Gopinath et al., 2021; Liu and Wang, 2021). Furthermore, due to the usual high-productivity requirements of manufacturing processes, some authors have researched the trade-off between productivity and safety (Scimmi et al., 2021; Zanchettin et al., 2019). Meanwhile, others have proposed adaptive robotic solutions for industrial applications (Levine and Williams, 2014, Levine and Williams, 2018; Villani et al., 2021). This large list of promising works has also come with some drawbacks. The lack of consensus on the meaning of concepts such as collaboration and adaptation has hindered a coherent development of methodologies and techniques. This has already shown to be a problem in safety applications, where the use of this terminology to assess risks might lead to confusion and potentially mistaken implementations (Vicentini, 2020).

A common approach to harmonize terminology and to enhance its reusability, is to use knowledge representation formalisms such as ontologies. Indeed, the use of ontologies has spread in the industrial domain, where the modular and reusable nature of this formalism has been of great help (Borgo et al., 2019; Karray et al., 2019; Mohd Ali et al., 2019). The 1872–2015 IEEE Standard Ontologies for Robotics and Automation (Schlenoff et al., 2012) presented a core ontology for robotics and automation, which is currently being extended to other robotics' sub-domains (Fiorini et al., 2017). Furthermore, ontologies have been widely used for autonomous

* Corresponding author.
E-mail address: aolivares@iri.upc.edu (A. Olivares-Alarcos).

https://doi.org/10.1016/j.compind.2022.103627
0166-3615/

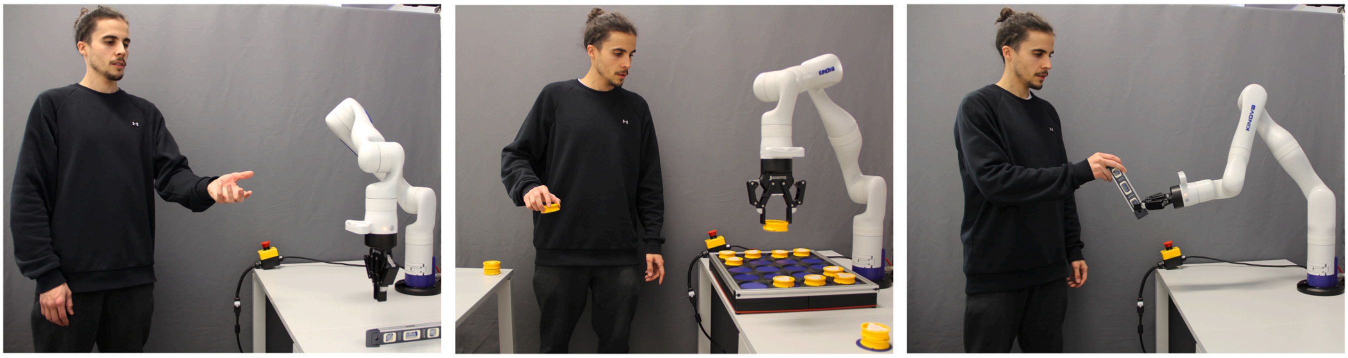

(a) (b) (c)

Fig. 1. Examples of collaborative tasks in which the human and the robot continuously share both the workspace and the execution of the task. (a) Example of a collaboration in which the human is stopped asking the robot for a tool. (b) Example of a kitting task in which the human and the robot collaboratively fill a tray with tokens. This task was used during the validation. (c) Example of a collaborative hand-over of a tool in which the robot and the human exchange forces.

robotics during the last years (Olivares-Alarcos et al., 2019), and we can even find some initial steps towards ontologies for collaborative robotics (Umbrico et al., 2020). However, in none of these works can we find a comprehensive analysis and formalization of the notions we are interested in: collaboration and adaptation.

In this article, we present the Ontology for Collaborative Robotics and Adaptation (OCRA). An ontology especially designed to represent the relevant knowledge in *collaborative scenarios* where robots *adapt* their plans to the ongoing changes in the environment. OCRA is the very first ontology that allows to formalize and reason about both human-robot collaboration and robot plan adaptation. Aiming to be inclusive and capture the common use of the relevant terminology in the literature, we studied other definitions before implementing our ontology. OCRA can be useful in different kinds of collaborative tasks, such as those shown in Fig. 1. Indeed, we selected one of those collaborative cases (see Fig. 1b) to validate OCRA's capability to answer a set of competency questions, and its robustness in some limit cases of the formalization. The article continues providing background on ontology in Section 2. OCRA is motivated and presented in Section 3, and its validation is discussed in Section 4 and Section 5. An analysis of the related work is shown in Section 6. Finally, we close the article with the conclusion and the future work, Section 7.

2. Background – Ontology

A first definition of ontology was given by Gruber (1993) stating: *an ontology is an explicit specification of a conceptualization*. Gruber's definition was informal and several authors tried to refine it either by specifying what a conceptualization is (Guarino and Giaretta, 1995), or by discussing further requirements like being *formal* and *shared* (Borst et al., 1997; Studer et al., 1998). Guarino et al. (2009) settled the issue proposing a formal definition which is today recognized in the community of applied ontology. An ontology is defined to be a *logical theory* consisting of a set of formulas whose models approximate as well as possible the intended models, i.e., those models that satisfy the conceptualization and the ontological commitments.

Since an ontology is a logical theory, it consists of individuals, classes, functions, relations and axioms. The exact list changes depending on the specific logic language one adopts. Usually, an ontology is given in First-Order Logic (FOL) (Smullyan, 1968), or in Web Ontology Language (OWL) (Antoniou and van Harmelen, 2004), both formal languages. *Individuals* are the objects in the ontology, the things the ontology is about. *Classes* are associated with properties and are used to identify the individuals that satisfy that property. *Functions* provide ways to identify and relate individuals. *Relations* are connections across individuals. *Axioms* are expressions in the language that use the previous elements to state what is true in the ontology.

2.1. Types of ontologies

Ontologies are usually divided in upper-level, reference, domain and application. An *upper-level ontology* is focused on general concepts like object, event, state or quality, and on high-level relations like parthood, constitution, participation or dependence. Examples are: SUMO (Suggested Upper Merged Ontology) (Niles and Pease, 2001), Cyc ontology (Lenat and Guha, 1990), and DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering) (Borgo et al., 2021). A *reference ontology* is an ontology that focuses on a discipline with the goal of fixing the general terms in it. It is highly reusable within the discipline, e.g., medical, engineering, enterprise, etc. (Guarino, 1998). When the ontology focuses on a more limited area, e.g., manufacturing or tourism, we call it a *domain ontology*. This kind of ontology provides vocabulary about the concepts within a domain and their relationships, about the activities taking place in that domain, and about the theories and elementary principles governing that domain. An *application ontology* contains all the definitions needed to model the knowledge required for a particular application, e.g., a Computer-Aided Design (CAD), a Computer-Aided Manufacturing (CAM) or an Enterprise Resource Planning (ERP) system.

A domain ontology, such as the one presented in this article, aims to make the domain knowledge explicit and formal, i.e., to fix, in a formal language, the vocabulary, its interpretation, and the valid assertions in the domain. A domain ontology can help to achieve reusability and interchangeability within and across communities (e.g., different industries) of both model and data. It is worth noting that there is no unique conceptualization of a domain.

2.2. Ontology languages and decidability

Some ontologies are expressed informally using ontology languages without formal semantics, e.g., Resource Description Framework (RDF) (McBride, 2004), and (part of) Unified Modeling Language (UML) (Booch et al., 2005). These languages are not suitable for automatic reasoning because there is no systematic way to constrain their interpretation. When reliability and reasoning are important, ontologies are expressed with formal languages, i.e., languages endowed with formal semantics like FOL and OWL. OWL is actually a family of languages including OWL DL and OWL Full. These languages are the most reliable and are of two types: decidable (e.g., OWL DL) and undecidable (e.g., FOL, OWL Full). An ontology that uses a decidable language is also called *computational* and it can be used at run-time for information extraction and verification. On the contrary, non-computational ontologies, such as FOL and OWL Full, although more expressive and precise in modeling knowledge, are not appropriate for reasoning at run time since they do not guarantee returning an answer to a query.

To take the most out of these languages, in this article we formalize the whole ontology using FOL. In this way we can express exactly what our notions mean. We also provide an OWL DL version of the ontology. It contains less knowledge but can be used for computational purposes and, therefore, implemented in a real robot for run-time reasoning.

3. OCRA – Ontology for collaborative robotics and adaptation

In this section we introduce OCRA, an Ontology for Collaborative Robotics and Adaptation. First, we define the scope and goal of our ontology, and we provide a set of competency questions. Second, we motivate the need for the different concepts we use in the ontology, we define them, and show their formalization in both FOL and OWL DL.

There are several methodologies to help the knowledge engineer in the ontology construction process, e.g., (Fernández-López et al., 1997; Spyns et al., 2008). Due to the variety of the possible cases and of the needed characteristics of the ontologies, none emerged as a definite standard. Furthermore, those methods are not suitable for developing an ontology from a foundational viewpoint where the characterization of the core concepts are more important than the coverage of the application domain, as in our case. Thus, our work relies on ontological analysis, an approach which precedes the usual ontology construction process and aims to fix the core framework for the domain ontology. This choice led us to perform the following steps: to set the ontology domain and scope (competency questions), to reconsider other conceptualizations (selection of relevant literature), to enumerate, analyse and compare existing concepts (identification of shortcomings), to develop and formalize a more solid conceptualization, and to create instances of the concepts and show their use (implementation/validation). Of course, there is some circularity in the actual procedure since this is a process of conceptual discovery and (re-)organization. As a final step, we also considered the documentation and maintenance of our proposal.

3.1. Scope, goal and competency questions

In order to develop OCRA, we followed a top-down approach. Hence, our ontology was built upon other higher level ontologies. Specifically, we developed an ontology that is compliant with Knowrob (Tenorth and Beetz, 2009; Beetz et al., 2018), the most widely used knowledge-based framework for robots. Therefore, we inherited the use of its upper ontology, the DOLCE+DnS Ultralite (DUL) foundational ontology?. Nevertheless, the concepts presented in this work are general enough to be adapted to and used with other upper ontologies. Furthermore, note that Knowrob is a system that has consistently improved along the last decade. This justifies using DUL and helps to generalize our work, since we could take advantage of some of their framework tools and experience.

OCRA was designed to represent relevant knowledge in the collaborative robotics domain, with special focus on collaboration and robot plan adaptation. We propose a group of questions, a set of requirements on the content, which scope and delimit the subject domain that has to be represented in the ontology. Particularly, we would like OCRA to be able to answer the following questions:

- Ontology coverage questions:
  - C1. What is a collaboration?
  - C2. What is a plan adaptation?
- Competency questions:
  - Q1. Which and how many collaborations are running now?
  - Q2. Which is the plan of a collaboration?
  - Q3. Which is the goal of a collaborative plan?
  - Q4. Are these agents collaborating?
  - Q5. Where is a collaboration happening?
  - Q6. How is a collaboration classified (e.g. non-physical)?
  - Q7. Which is the risk of a collaboration?
  - Q8. Which and how many plan adaptations are running now?
  - Q9. Which is/are the agent/s participating in the plan adaptation?
  - Q10. Why is an adaptation of an agent's plan happening?
  - Q11. Which is the plan before and after an adaptation?
  - Q12. Which is the goal of the agent involved in the adaptation that is also the goal to be achieved by both the old and the new plan?

3.2. On the meaning of collaboration

In this section, we discuss the need for a formal definition of `Collaboration`, and present our definition and formalization.

3.2.1. Rationale – Ambiguity in the literature

The Oxford Dictionary defines `Collaboration` as *'the act of working with another person or group of people to create or produce something'* (Oxford-University, 2022). This informal definition would let us talk about collaborative events. However, a formal definition is needed to enable robots to reason about these events. In this section, we analyze several informal definitions from the literature, highlighting their differences and common points, and motivating the need for a comprehensive formal model for `Collaboration`.

In 2011, the International Organization for Standardization released the ISO 10218.1 (ISO, 2011b) and the ISO 10218.2 (ISO, 2011a), which defined `Collaboration` as *'a special kind of operation between a person and a robot sharing a common workspace'*. Vicentini (2020), discussed the ambiguity in the collaborative robotics' terminology. He stated that at least, *'there is a predominant consensus in assigning the concept collaboration to continuous, purposeful interaction associated with potential or accidental physical events (contacts)'*. The Organization for Economic Co-operation and Development, defined the collaborative problem-solving competency as: *'the capacity to engage in a process whereby two or more agents attempt to solve a problem by sharing the understanding and effort required to come to a solution'* (O. for Economic Co-operation, 2017). Oliveira et al. (2007), defined collaboration session (CS) as *'an event that is composed of the actions of its participants. A CS has one or more objectives, defining its main purpose'*. Dillenbourg, 1999, discussed about the definition of collaborative learning. He stated that *'collaborative situations involve symmetry between what agents know and do, shared goals, and a low division of labor'*. Silverman (1992) defined `Collaboration` as *'the mutual sharing of goals in completing the tasks'*. Terveen (1995), defined `Collaboration` as *'a process in which two or more agents work together to achieve shared goals'*. He also derived a set of fundamental issues from his definition: agreement on the goal, plan and coordination, shared context and understanding of the current situation, communication, and adaptation and learning. Kolfschoten (2007) studied several definitions of `Collaboration` and proposed a refined one: *'a joint effort toward a goal. This implies that all participants make effort, combine it and direct it to achieve a desired state or outcome (goal)'*. Bauer et al. (2008), surveyed the human-robot collaborative domain, for them, collaboration means *'working with someone on something, aiming at reaching a common goal. To work cooperatively on something the partners need to agree on a common goal and a joint intention (plan) to reach that goal'*. Ajoudani et al. (2018), reviewed the state-of-the-art on human-robot collaboration. They considered that human-robot collaboration *'falls within the*

Table 1
Set of main aspects related to 'Collaboration' extracted from the literature. 'Formal' shows whether the literature definition was formalized or not. 'Goal', 'Plan' and 'Interaction/Execution' columns indicate whether the notion of each aspect was captured or not by the definition. *Implicit in the definition.

Ref.	Formal	Goal	Plan	Interaction/ Execution
(ISO, 2011a)	No	–	–	Yes
(Vicentini, 2020)	No	Yes	–	Yes
(O. for Economic Co-operation, 2017)	No	Yes	Yes	Yes
(Oliveira et al., 2007)	Yes	Yes	–	Yes
(Dillenbourg, 1999)	No	Yes	Yes	Yes
(Silverman, 1992)	No	Yes	–	–
(Terveen, 1995)	No	Yes	Yes	Yes
(Kolfschoten, 2007)	No	Yes	Yes	Yes
(Bauer et al., 2008)	No	Yes	Yes*	–
(Ajoudani et al., 2018)	No	Yes*	–	Yes
(Umbrico et al., 2020)	Yes	Yes*	Yes*	–
Ours	Yes	Yes	Yes	Yes

general scope of human-robot interaction, and it is defined when human(s), robot(s) and the environment come to contact with each other and form a tightly coupled dynamic system to accomplish a task'. Note that the interaction or contact might also be non-physical (e.g., mental). Umbrico et al. (2020), defined the concept collaborative process as *'a process, in order to represent production events that modify over time the state of the production environment from an initial situation to a final/resulting one'*. Their formal definition was the closest one to ours, although it lacked explicit mention to the shared plan and goal, focusing too much on the situation's change. Hence, we think that our definition is more general, and this might be considered as a specialization of ours.

Even though all these definitions diverge, it is possible to find some patterns that most of them follow: collaborative agents must share a goal and a plan (understanding/coordination), and there must be interaction (Borgo, 2019) between them while executing the plan. Table 1 depicts a summary with these main aspects of `Collaboration` for each of the studied articles.

3.2.2. Definition in natural language

Considering all the aforementioned definitions, `Collaboration` is usually defined as a special kind of spatio-temporal entity (an event). Furthermore, it is often related to a goal and a plan, and it requires interaction among the agents. Hence, we define `Collaboration` as *'an event in which two or more agents share a goal and a plan to achieve the goal, and execute the plan while interacting'*. We use interaction as an unspecified term as it belongs to a higher level ontology, it is thus considered a primitive concept in OCRA. Informally speaking, interaction is *'the act of communicating with somebody, or having an effect on each other'* (P. Oxford-University, 2022). For example, during a collaboration, a robot and a human interact when they exchange forces, and also when the robot is sharing its perception about the safety situation (e.g., by voice or lights). Note that our definition states that the collaborative agents shall share a plan and the goal to be achieved. Hence, even when an agent delegates a part of a plan, we understand that the agent maintains co-responsibility for that part of the plan.

For example, let's consider that there is a robot and a human that are collaborating to fill the different compartments of a tray with work pieces. We will say that the collaboration exists as long as:

- the robot and the human share a Plan to fill the tray. Note that the plan can include generic activities, like 'picking some pieces and placing them on the tray until it is full', or more specific activities like 'picking the pieces starting from the closest and placing them on the tray from left to right';
- the robot and the human share the Goal to be achieved by the execution of the shared plan. It may be general, e.g., 'all tray's compartments with a piece', or specific, e.g., 'each tray's compartment filled with a certain piece'; and
- the robot and the human execute the shared plan while they interact, thus, they share an understanding of who is in charge of what during the execution.

3.2.3. Formalization in FOL

Formalizing our definition of `Collaboration`, we reused as many classes and relationships as possible from the foundational ontology DUL (note the prefix *'dul.'*). The final formalization in FOL is:

$$
\begin{aligned}
&Collaboration(e) \equiv dul.Event(e) \wedge \\
&\quad \exists\, y, z, p, g, t\ (y \neq z) \wedge \\
&\quad \mathrm{dul}.Agent(y) \wedge dul.hasParticipant(e, y) \wedge \\
&\quad \mathrm{dul}.Agent(z) \wedge dul.hasParticipant(e, z) \wedge \\
&\quad \mathrm{dul}.Plan(p) \wedge dul.Goal(g) \wedge \\
&\quad \mathrm{dul}.hasComponent(p,g) \wedge executesPlan(e,p) \wedge \\
&\quad \mathrm{dul}.hasTimeInterval(e,t) \wedge \\
&\quad \forall\, x\ (dul.Agent(x) \wedge dul.hasParticipant(e,x)) \rightarrow \\
&\quad \mathrm{hasPlan}(x,p,t) \wedge hasGoal(x,g,t).
\end{aligned} \tag{1}
$$

The definition reads as follows: *a collaboration is an event (e) in which at least two agents (y and z) participate, it is the execution of a plan (p) with some goal (g), and for any agent (x) in the collaboration its aim is to execute that plan and to achieve that goal.*

Note that we did not restrict the definition stating that the pursued goal must be achieved at the end of the collaboration, thus, being general and considering cases in which the goal of the collaboration is not achieved (and perhaps, unknown to the agents, even not achievable). Furthermore, the relationship `executes plan` was used here as a primitive which means 'following the sequence of actions in the plan'. Hence, we did not consider this notion in the strictest sense, which would be to execute the whole plan. This predicate holds between an event and a plan that is executed by that event. We also found necessary the use of two new relationships that were not explicitly defined in DUL: `has plan` and `has goal`. They relate an agent with a plan and a goal, respectively, during a time interval. First, `has plan` means that 'an agent intends to execute a sequence of actions (plan)'. Second, `has goal` implies that 'an agent desires to achieve a goal'.

3.3. On the meaning of adaptation

In this section, we discuss the need for a formal definition of `Adaptation`, and present our definition and formalization.

3.3.1. Rationale – Ambiguity in the literature

The Oxford Dictionary P. Oxford-University (2022) defines `Adaptation` as *'the action or process of changing something, or of being changed, to suit a new purpose or situation'*. This informal definition would be helpful to talk about adaptation events. However, a formal definition is needed to allow robots to reason about these events. In this section, we analyze several informal definitions from the literature, spotlighting their discrepancies and shared points, and encouraging the need for a comprehensive formal model for `Adaptation`.

Järvenpää et al. (2016), presented an adaptation approach for small-size production systems, in which `Adaptation`*'referred to all controlled changes the production system goes through during its life cycle'*. Martín H et al., 2009, proposed a mathematical model of the phenomenon of `Adaptation`. Specifically, they defined a Law of Adaptation: *'every adaptive system converges to a state in which all* *kind of stimulation ceases'*. For them, an adaptive system *'has at least one process which controls the system's adaptation to increase its efficiency to achieve its goals'*. Lints (2010), identified and discussed the main aspects of adaptation from different fields of research. He defined `Adaptation` as *'a process to change something (itself, others, the environment) so that it would be more suitable or fit for some purpose than it would have been otherwise'*. Smit and Wandel (2006), reviewed the concept of adaptation regarding humans' adaptation to global changes such as climate change. The authors stated that `Adaptation`*'might refer to a process, action or outcome in a system, in order for the system to better cope with, manage or adjust to some changing condition, stress, hazard, risk or opportunity'*. Smit et al. (2000), discussed that a thorough description of adaptation should specify the system of interest that adapts, the stimulus that causes the adaptation, and the involved processes and their outcomes. Gjorven et al. (2006), considered `Adaptation` as a service, and defined it as *'a service whose input event is an adaptation trigger, and whose output events are a set of services that potentially has been modified or produced during the adaptation'*. All these definitions are ambiguous, but there are some patterns that most of them follow: adaptation shall be triggered by a stimulus, shall occur on an entity that would change to a new state, and shall aim to continuously pursue the achievement of a goal. Table 2 depicts a summary with these main aspects of `Adaptation` for each of the definitions.

3.3.2. Definition in natural language

After studying the state of the art, we realized that providing a general definition of `Adaptation` would be extremely challenging. Barandiaran et al. (2009), discussed that adaptation involves a norm specifying which is the appropriate change to make. Hence, depending on the type of norm, we could find different types of adaptations: task or plan-based, evolutionary, ecological, etc. In our work, we focused on plan-based adaptations, changes aimed at continuously pursuing the completion of a goal given an unexpected state or situation.

For us, `Plan Adaptation` is *'an event in which one (or more) agent, due to its evaluation of the current or expected future state, changes its current plan while executing it, into a new plan, in order to continuously pursue the achievement of the plan's goal.'*.

From our definition, we can extract the conclusion that if a plan was changed before starting its execution, that would not be an adaptation. Also note that if a change was part of a plan, we would not consider it to be an adaptation. Hence, if a robot's plan included two optional executions, choosing one would not be an adaptation. Indeed, some authors claimed that the capacity to adapt depends on the observer who chooses the scale and granularity of description (Paolo, 2005; Barandiaran and Moreno, 2008). For instance, in a micro-scale, obstacle avoidance might be seen as an adaptive behavior, but in an environment rich in obstacles, it would not.

For instance, let's consider the previous example where a robot and a human collaborate to fill the different compartments of a tray with work pieces. We will say that an adaptation exists as long as: .

Table 2
Set of main aspects related to 'Adaptation' extracted from the literature. 'Formal' column shows whether the literature definition was formalized or not. 'Trg.' (trigger), 'Ent.' (entity), 'Chg.' (change) and 'Goal' columns indicate whether the notion of each aspect was captured or not by the definition. *Implicit in the definition. **Mathematical model but not an ontological one.

Ref.	Formal	Trg.	Ent.	Chg.	Goal
(Järvenpää et al., 2016)	No	–	Yes	Yes	–
(Martín H et al., 2009)	Yes**	Yes	Yes	Yes	Yes
(Lints, 2010)	No	–	Yes	Yes	Yes
(Smit and Wandel, 2006)	No	Yes	Yes	Yes*	Yes*
(Smit et al., 2000)	No	Yes	Yes	Yes*	–
(Gjorven et al., 2006)	No	Yes	Yes	Yes*	–
Ours	Yes	Yes	Yes	Yes	Yes

- the robot has a plan, and it is executing it while the perception of a current or future state (situation) triggers the adaptation. A possible plan could be 'moving to a compartment to release a piece', and the trigger might be 'the compartment is full'; and
- the robot changes its plan by no longer executing the action required by the previous plan, and from now on executes the new plan. Still aiming to fill the tray, the new plan could be 'moving to another free compartment'.

3.3.3. Formalization in FOL

In order to formalize our definition of `Adaptation`, again we reused as many classes and relationships as possible from the foundational ontology DUL (note the prefix *'dul.'*). The final formalization in FOL is:

$$\begin{aligned}
&PlanAdaptation(e) \equiv dul.Event(e) \wedge \\
&\exists\, s, g, a, o, n, i, f, p, q\ dul.Situation(s) \wedge dul.Goal(g) \wedge \\
&\quad \text{dul.}Agent(a) \wedge \text{dul.}hasParticipant(e, a) \wedge \\
&\quad \text{dul. } Plan(o) \wedge dul.\ hasComponent(o, g) \wedge \\
&\quad \text{dul. } Plan(n) \wedge dul.\ hasComponent(n, g) \wedge \\
&\quad \text{dul. } hasPostcondition(i, s) \wedge betterPlan(s, n, o) \wedge \\
&\quad \text{dul. } Event(i) \wedge dul.\ hasTimeInterval(i, p) \wedge \\
&\quad \text{dul. } Event(f) \wedge dul.\ hasTimeInterval(f, q) \wedge \\
&\quad p < q \wedge i + f = e \wedge executesPlan(i, o) \wedge \\
&\quad \text{executesPlan}(f, n) \wedge \neg\, executesPlan(f, o) \wedge \\
&\quad \forall\, x\ ((dul.\ Agent(x) \wedge dul.\ hasParticipant(i, x)) \rightarrow \\
&\qquad \text{hasPlan}(x, o, p) \wedge hasGoal(x, g, p)) \wedge \\
&\quad \forall\, x\ ((dul.\ Agent(x) \wedge dul.\ hasParticipant(f, x)) \rightarrow \\
&\qquad \text{hasPlan}(x, n, q) \wedge hasGoal(x, g, q)).
\end{aligned} \tag{2}$$

The definition reads as follows: *a plan adaptation is an event (e) with at least one agent (a), which is the change of a plan (o) with a goal (g) into a new plan (n) with the same goal, where the change is due to the evaluation that the situation (s) holding after the first part of the event (i) makes plan (n) in the second part (f) better than continuing plan (o), and in the first part any agent (x) aims to execute plan (o), while in the second part any agent aims to execute the plan (n), and every agent has always the same goal (g).*

Note that at least one agent participates in the whole adaptation event while other agents may change due to the adaptation. We included a new predicate/relationship: `better plan`, which relates two plans and a situation that makes one of the plans better to achieve a goal. Hence, one could use this relation to state that a situation has caused that one plan is no longer good, and a new plan is better to accomplish a goal. Note that we could similarly define `worse plan` as its inverse predicate.

3.4. Complementary terminology

In this section, we address other relevant terms of the industrial collaborative robotic domain. We first discuss the need for the chosen set of terms, and then we define and formalize them as part of OCRA.

3.4.1. Rationale – Common terms in the literature

In the collaborative robotics literature, apart from the concept of `Collaboration`, it is widely spread the use of terms such as workspace, safety, or collaboration types (Vicentini, 2020; ISO, 2011b,a). Most of this terminology is already defined in well established ISO standards, so we could just reuse the notions. However, there are no formal standard definitions yet, so we formally defined the concepts as part of OCRA. Hence, allowing robots to reason about the place where they collaborate, safety aspects, and the different types of collaboration. The ISO 10218.1 ISO (2011b) defined

Collaborative Workspace as *'a workspace within the safeguarded space where the robot and a human can perform tasks simultaneously during production operation'*. This definition is broad enough to capture most of the collaborative scenarios found in the industry, where a fixed workspace is often designed for collaborations. However, aiming to be general, we also considered this concept from the perspective of the place/environment where a collaboration happens. In that case, the place could dynamically change due to the collaboration needs (e.g., if the collaborators have to do operations using machines in different areas of the shop floor). Indeed, some authors defined a collaborative dynamic geometrical region that includes the intersection of both the robot's and the human's workspaces (Melchiorre et al., 2021). Hence, we decided to define two different concepts: one for the notion of the place where a collaboration occurs (Collaboration Place), and another one for the common industrial notion of a fixed place for collaborations (Collaborative Place). Regarding safety, the standard is to follow the guidelines of the ISO 12100 (ISO, 2010), which focused on machinery's risk assessment and risk reduction. ISO 12100 defined risk as *'combination of the probability of occurrence of harm and the severity of that harm'*. In this work, we formalized and defined the concept Collaboration Risk (see Section. 3.4.2 for more details).

Finally, it would be interesting to classify different types of collaboration, so that robots could behave differently depending on each type. The ISO 10218.2 (ISO, 2011a) defined four different collaborative operational modes for robots. They are useful to talk about different robot behaviors or strategies, but they cannot directly be considered as sub-classes of collaboration. Bauer et al., 2016, proposed a classification of different collaboration levels: cell, coexistence, synchronized cooperation, and collaboration. However, these categories are ambiguously used in the literature (Vicentini, 2020), and they might lead to confusion. For the time being, there is not a standard taxonomy of collaboration types, actually, there can be many depending on the application domain. Hence, we focused on a classification that is relevant for our application and for risk analysis, which is based on the degree of physical human-robot interaction: Non-physical Collaboration, Indirectly Physical Collaboration, and Directly Physical Collaboration. In the future, other classifications might be considered to extend OCRA.

3.4.2. Definition and formalization

In OCRA, a Collaboration Place*'is the spatial location or the place of a collaboration'*. This concept was formally defined as a subclass of Place (DUL) that is location of a Collaboration. Note that this definition focuses on the existence of a collaboration and where it is located. Hence, a collaboration place is the union of the spatial locations of all the entities involved in the collaboration, which could change across time. For instance, if the agents involved in the collaboration move to other places, the collaboration place would also move.

We defined Collaborative Place as *'a role of a place that is specifically dedicated to collaborations'*. It was formalized as a subclass of Role (DUL) that classifies a Place. Recall that this definition focuses on the place where collaborations can occur. It is meant to capture the traditional view of industrial collaborative workspaces, in which a collaboration is only considered inside of a fixed workspace. It is worth noting that when a collaboration occurs in a place whose role is to be a Collaborative Place, there is also a Collaboration Place that can be different to the first one. For instance, when a collaboration is occurring at a work cell that plays the role of a Collaborative Place, if one of the agents goes out of the work cell to do part of the collaboration, the place where the collaboration happens (the Collaboration Place) would be different to the work cell.

Concerning safety, and based on the ISO 12100 (ISO, 2010), Collaboration Risk was defined as *'a quality that has a value used to characterize a collaboration, or a part of it, which combines the probability of occurrence of a given harm and the severity of that harm during that collaboration'*. It was formalized as a sub-class of Quality (DUL) that is quality of a Collaboration.

Regarding the different types of collaboration, we first defined Non-physical Collaboration as *'an event type that classifies a collaboration, or a part of it, in which the involved agents do not exercise any physical force'*. For instance, selecting the next part of the plan to execute, asking for a tool (see Fig. 1a), verbally communicating commands or recommendations to collaborators (Chacón et al., 2020; Nikolaidis et al., 2018), or monitoring how a part of a plan is executed by another agent. Second, Indirectly Physical Collaboration is *'an event type that classifies a collaboration, or a part of it, in which the involved agents exercise physical forces but they do not physically restrict the freedom of movement of any of the other agents'*. For instance, when a robot moves close to a moving human without exchanging forces (see Fig. 1b). Third, Directly Physical Collaboration is *'an event type that classifies a collaboration, or a part of it, in which the involved agents exercise physical forces, and they do physically restrict to some degree the freedom of movement of at least one of the agents'*. This includes the cases where the involved agents exchange contact forces, directly or through an object, as shown in Fig. 1c. Hence, any movement of one of the agents would affect some other agent. For instance, a collaborative hand-over (Pan et al., 2019), the collaborative task of polishing an object (Olivares-Alarcos et al., 2019), or the assembly of a piece of furniture (Rozo et al., 2013). In the case in which the freedom of movement is restricted by rules such as those related to safety (e.g., robot stops if human is closer than a given distance), a collaboration would still be considered as Indirectly Physical Collaboration because the movement of the robot is restricted by the safety behavior, not by a pure physical impediment. Of special interest would be the case when a robot and a human (Pan et al., 2019), or two collaborative robots (Garcia-Camacho et al., 2020), are holding a deformable object. If both agents held it close enough so that there was still freedom of movement, we would consider it as a Indirectly Physical Collaboration. If they went further, so that the deformable object is completely stretched/extended, then we would be in a Directly Physical Collaboration, because if one of the agents moved some of the others would be affected by it. This example is related to a collaborative hand-over, a task that might also be categorized under other characteristics (e.g., robots' adaptability/responsiveness). In the future, we will consider other features and tasks to enlarge the list of collaboration types included in OCRA. The three concepts were defined as sub-classes of Event Type (DUL) that classify a Collaboration. We considered the option of defining them as sub-classes of Collaboration. Nevertheless, although they were useful concepts for reasoning, the differences between them were not ontologically meaningful for us. Furthermore, note that we intentionally avoided the use of the concept 'contact' in our definitions. We think that our classification is more general since we also consider other physical forces, not only whether or not the human and the robot are touching each other.

3.5. OCRA formalization in OWL

Complementing the formalization in FOL, we also provide an OWL DL version of the ontology. It contains less knowledge but can be used for computational purposes and can be implemented in the robot for run-time reasoning. The ontology was implemented using Protégé (Gennari et al., 2003), and we publish the developed OWL file together with other additional material[1] to facilitate reuse and comparison. Most of the axioms we defined in FOL were translated

[1] www.iri.upc.edu/groups/perception/OCRA

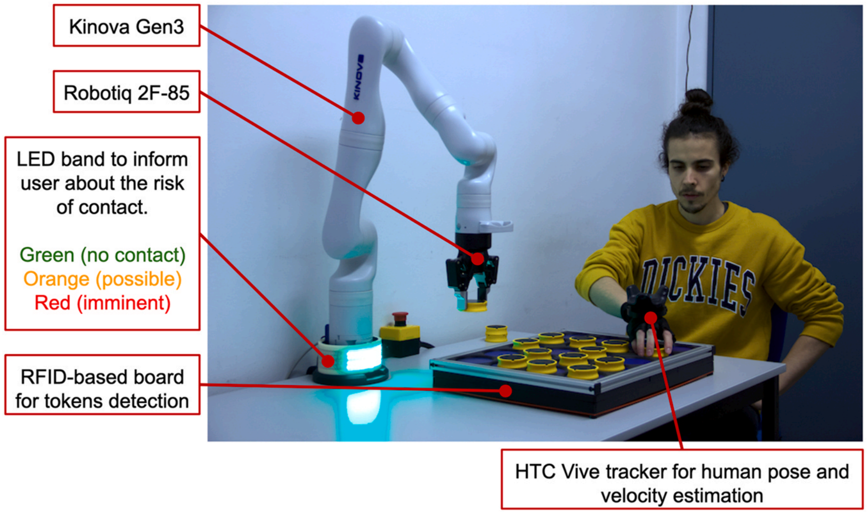

Fig. 2. Setup of collaboratively filling a tray: example of an industrial kitting task used during the validation of this work.

into OWL DL, with the exception of the three ternary relationships: `has plan`, `has goal`, and `better plan`. FOL supports the use of ternary relationships, but OWL DL does not (although in some cases one can overcome this problem (Rector and Noy, 2006). First, `has plan` and `has goal` were defined as ternary to express that agents had a goal or a plan during an interval of time. However, in the OWL DL version of OCRA, the two properties only relate agents with their plans and goals, without stating for how long those relationships hold. This is not necessarily critical since the use of OWL DL at run-time happens while the agents do have the plan and goal. For a broader use of the OWL DL formalization, this issue can be solved by 'reification', introducing several relations *hasGoal_tp*(r, g), one for each instant (tp) in which the relation holds. For instance, if the plan is to move objects and we do that at frequency 1 per minute for a total of 1 h, one could imagine to check every minute whether the agents maintain the plan. For this it is enough to introduce 3600 relationships *hasGoal_tn*(r, g) with n going from 0 (initial time) to 3599. This solution is activity dependent so we do not present it in the general definition. Furthermore, one can also exploit a temporal history of the knowledge base's facts (episodic memories) (Beetz et al., 2018). Hence, we could determine the temporal interval during which a relationship holds (e.g., the time an agent has a goal or a plan). Second, `better plan` was defined as ternary to model that, given a situation, a plan is better than another plan. In the OWL DL version of OCRA, there is one relationship that substitutes the ternary one: `is better plan than`, relating two plans. Since the relationship `is better plan than` is evaluated at the time when the situation s holds, the notion formalized in OWL DL is a good approximation of the original one in FOL. Finally, note that we also include other complementary relationships (e.g., the inverse of all the previous ones).

4. Validation I – Answering the competency questions

In this section, we present a qualitative validation of the use of OCRA in a lab mock-up of a real task, where a robot and a human share the task of filling the compartments of a tray (see Fig. 2). The video of one of the experiments can be found in the additional material.[2] This validation was meant to evaluate the ontology's capabilities to answer the set of competency questions proposed in Section 3.1. Note that the design requirements of the ontology were those competency questions. Hence, answering them proves that the ontology was properly formalized, and that it meets the desired prerequisites. Specifically, we contextualized the competency questions showing different situations extracted from the proposed collaborative scenario. For each situation, we used an OWL DL knowledge base populated with the proper instances to answer the queries. Note that since we are able to reason over OCRA using an inference engine (HermiT (Glimm et al., 2014), we also validate that the ontology is consistent and coherent. In order to reliably compute the risk of human-robot collision, we used the pose and velocity extracted from an HTC Vive tracker attached to the human's hand. Inspired by Hou et al. (2014), we utilized the velocity and pose of both the human and the robot to compute the Time-To-Contact (TTC). TTC is the time that would take the robot's end effector and the human's hand to collide if they kept moving at the same relative velocity. Note that the measurements from the HTC and the robot were taken at 100 Hz. When the TTC was lower than a certain threshold, the robot stopped (high risk of collision). The medium degree of risk corresponded to when TTC was greater than the threshold and different to infinite. When TTC was infinite, meaning that there was no expected contact, the level of risk was low. Using the lights on the robot's base, the robot shared its interpretation of the collision's risk with the operator (Fig. 2). An RFID-based board was used for a fast and precise token-compartment detection. We ran all the software in a desktop PC with an Intel Core i7–7800X CPU (12x 3.50 GHz), a 32 GB DDR4 RAM, and an NVIDIA GeForce FTX 1080 Ti/PCLe/SSE2 GPU.

4.1. Filling a tray – Application ontology

In order to represent the knowledge of the proposed use case, we needed some extra concepts. They were defined as either instances or specializations of DUL's classes. In the scenario of filling a tray, there were different objects: the robot, the human operator, the board (tray), the compartments, and the tokens. All of them were instances of different sub-classes of `PhysicalObject` in DUL: *'any Object (DUL) that has a proper space region'*. We defined the robot and the human as instances of `PhysicalAgent`, and the board, the compartments, and the tokens as instances of `DesignedArtifact`. For the board and the compartments, we included a new class: `AvailableCapacity`, defined as a `Quality` in DUL. This quality lets us capture the knowledge about whether a compartment or the tray is already filled or not.

[2] www.iri.upc.edu/groups/perception/OCRA

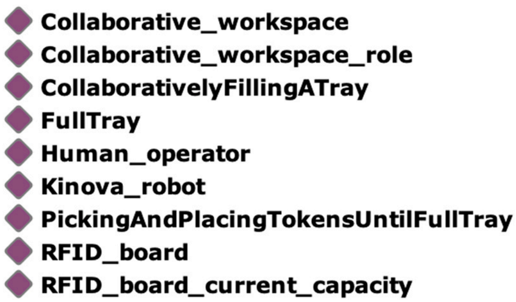

Fig. 3. ABox to answer general competency questions about collaboration in Protege.

4.2. Part 1 – Questions about collaboration

In this section, we show how OCRA is able to answer general questions about collaboration. Specifically, we address here the first coverage question and the first five competency questions presented in Section 3.1. Once we imported DUL and OCRA, we instantiated in our knowledge base all the entities that were involved in the collaboration (see Fig. 3).

First, we can query the knowledge base to determine what is a collaboration (C1) and which are its current instances (Q1):

1. What is a collaboration? Which and how many collaborations are running now? The result shows that a collaboration is an event and there is one existent collaboration: collaboratively filling a tray.

```
DL query:
  Collaboration
Query results:
  Direct superclasses:
    Event
  Instances:
    CollaborativelyFillingATray
```

Once we know that there is a collaboration, we can use it to further query the knowledge base for the plan executed in the collaboration (Q2):

2. Which is the plan of a collaboration? The plan of the current collaboration is to pick and place tokens on the tray until it is full.

```
DL query:
  isPlanExecutedIn value CollaborativelyFillingATray
Query results:
  Instances:
    PickingAndPlacingTokensUntilFullTray
```

Using OCRA, we can also ask for the goal of the plan executed in a collaboration (Q3):

3. Which is the goal of a collaborative plan? The goal of the plan executed in the current collaboration is to fill the tray with tokens.

```
DL query:
  isComponentOf some (Plan and isPlanExecutedIn value
    CollaborativelyFillingATray)
Query results:
  Instances:
    FullTray
```

To know which are the agents that are collaborating to fill the tray, we can ask the next query to the knowledge base (Q4):

4. Are these agents collaborating? The two agents which collaborate are the Kinova robot and the human operator.

```
DL query:
  isParticipantIn some Collaboration and hasGoal value
    FullTray and hasPlan value
    PickingAndPlacingTokensUntilFullTray
Query results:
  Instances:
    Kinova_robot
    Human_operator
```

Finally, it is possible to know at which place the collaboration is happening (Q5):

5. Where is a collaboration happening? The current collaboration is happening at the Collaboration Place 'Collaborative workspace' that is also playing the role of a 'Collaborative Place'.

```
DL query:
  isLocationOf value CollaborativelyFillingATray and
    isClassifiedBy value
    Collaborative_workspace_role
Query results:
  Instances:
    Collaborative_workspace
```

4.3. Part 2 – Questions about collaboration types and risk

In this case, we tackled the next two competency questions, those about the type and the risk of a collaboration. Our knowledge base contained the same content as before plus some instances about the specific type of collaboration and the risk (see Fig. 4).

From the same collaboration event, we extracted different situations for each collaboration type and risk. We considered three different types of collaboration corresponding to the ones defined in Section 3.4. In Fig. 5, we depict a picture for each of the types. We focused on the risk of collision, and we defined three different levels: high, medium and low (see Fig. 6).

We start querying the type of the current collaboration (Q6). Note that there can be three different answers for the same query, depending on the cases shown in Fig. 5. To avoid repetition, we only show the case in which the collaboration was classified as `Non-Physical`:

6. How is a collaboration classified? The current type of collaboration is non-physical.

```
DL query:
  classifies value CollaborativelyFillingATray
Query results:
  Instances:
    CollaborativelyFillingATray_non-physical
```

The last query of this section is about the risk of a collaboration (Q7), which changes over time. However, using the OWL DL reasoners from Protégé we would always get the same answer. Because they only work with classes and instances, and we stored the specific value of risk as the data value of an entity. In order to overcome this, and to avoid introducing other languages such as SPARQL (Pérez et al., 2006), SQWRL (O'Connor and Das, 2009), we restricted the data value of the current risk in the query. Hence, if the entity containing the current risk had the queried data value, it would be returned as a result. Otherwise, the result would be empty. In this case, the risk in the ABox was set to 'high':

7. Which is the risk of a collaboration? The current risk of the collaboration is 'high' because the result of the query is not empty.

```
DL query:
  CollaborationRisk and isQualityOf value
    CollaborativelyFillingATray and
    hasDataValue value 'HIGH_RISK'
Query results:
  Instances:
    CollaborativelyFillingATray-current_risk
```

4.4. Part 3 – Questions about adaptation

In this section, we address the rest of the coverage and competency questions, which are about plan adaptation. Once we imported DUL and OCRA, we instantiated in our knowledge base all the entities that were involved in the adaptation event (see Fig. 7).

We propose a new situation for the competency questions about adaptation. The robot modified its symbolic task plan and continued with the other free targets after the human filled one of the compartments the robot was meant to fill (see Fig. 8). Note that the robot's path planning was a simple point to point straight navigation. Aiming to show the reusability of OCRA, we also considered another adaptation triggered by an imminent human-robot collision. The robot adapted its plan by stopping and remaining in admittance mode until the human commanded it to resume the motion. However, to save space in the article, we explained it in the additional material.[3]

We start with general questions about what is a plan adaptation (C2) and which are the current instances (Q8):

8. What is a plan adaptation? Which and how many plan adaptations are running now? The result shows that a Plan Adaptation is an Event and there is one existent adaptation: full compartment-based adaptation

```
DL query:
  PlanAdaptation
Query results:
  Direct superclasses:
    Event
  Instances:
    Full_compartment_adaptation
```

Once we know the current plan adaptation, we can ask for the agent/s involved in it (Q9):

9. Which is/are the agent/s participating in the plan adaptation? The result returns that the agent is the Kinova robot.

```
DL query:
  isParticipantIn value Full_compartment_adaptation
Query results:
  Instances:
    Kinova_robot
```

The trigger or the cause of the plan adaptation can be retrieved using the following query (Q10):

10. Why is an adaptation of an agent's plan happening? The result shows that the stimulus of the adaptation was that the target compartment (19) was full.

```
DL query:
  Situation and isPostconditionOf some (Event
    and isPartof value
    Full_compartment_adaptation and
    executesPlan some (Plan and
    isWorsePlanThan some Plan))
Query results:
  Instances:
    RFID_board_compartment_19_is_full
```

We could further ask for the details of the adaptation's cause, which are also represented using OCRA. The current situation is setting of the compartment nineteen, whose available capacity is zero. This indicates that the compartment is full, the reason why the robot adapts its plan. We could check that this is true querying the knowledge base:

11. Details of the adaptation's cause. The result shows that the compartment nineteen was setting for the situation and that its available capacity was zero.

```
DL query:
  hasSetting value
    RFID_board_compartment_19_is_full and
    hasQuality some (AvailableCapacity and
    hasDataValue value 0)
Query results:
  Instances:
    RFID_board_compartment_19
```

Using OCRA we can also ask for the plan before and after the adaptation (Q11):

12. Which is the plan before the adaptation? The plan was the initial one: filling compartment 19.

```
DL query:
  isPlanExecutedIn some (Event and  isPartOf value
    Full_compartment_adaptation) and
    isWorsePlanThan some Plan
Query results:
  Instances:
    Full_compartment_adaptation_initial_plan
```

13. Which is the plan after the adaptation? The plan was the final one: filling compartment 4.

```
DL query:
  isPlanExecutedIn some (Event and  isPartOf value
    Full_compartment_adaptation) and
    isBetterPlanThan some Plan
Query results:
  Instances:
    Full_compartment_adaptation_final_plan
```

Finally, we can query the knowledge base to get the goal of the plans that is also the goal of the agent/s involved in the adaptation (Q12):

14. Which is the goal of the agent involved in the adaptation that is also the goal to be achieved by both the old and the new plan? The goal is to fill the whole tray.

```
DL query:
  isComponentOf value
    Full_compartment_adaptation_initial_plan and
    isComponentOf value
    Full_compartment_adaptation_final_plan and
    isGoalOf some (Agent and isParticipantIn value
    Full_compartment_adaptation)
Query results:
  Instances:
    FullTray
```

5. Validation II – Limit cases evaluation

In this second validation, we study the robustness of the proposed ontological model, analysing OCRA's performance in several limit cases of the formalization. Particularly, we propose a set of examples of `Collaboration` and `Plan Adaptation` that contain incongruent or incomplete axioms. We explore how the formal definitions in FOL and OWL DL behave in these cases, observing whether OCRA is able to exclude or not the incorrect instances. Table 3 depicts the description of the proposed cases, the

[3] www.iri.upc.edu/groups/perception/OCRA

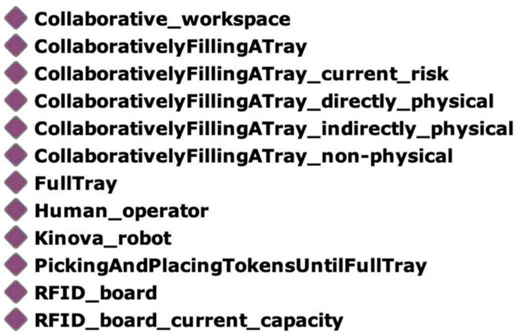

Fig. 4. ABox to answer competency questions about the type and the risk of a collaboration in Protege.

classification into correct instances or not, and the relevant axioms in FOL and OWL DL. The results show that the formal definitions within OCRA allow to exclude those cases in most of the situations. This validation proves the strength of our formal model in situations where it might be unclear whether an event is or not a `Collaboration` or a `Plan Adaptation`.

6. Related work

The 1872–2015 IEEE Standard Ontologies for Robotics and Automation (Schlenoff et al., 2012) was conceived as a reference for knowledge representation and reasoning in the domain, and a formal vocabulary for humans and robots to share knowledge about robotics and automation. However, it did not cover terminology for particular robotic sub-domains. Hence, several ontology-based systems for autonomous robots were implemented focusing on more specific notions. Some examples are Knowrob (Tenorth and Beetz, 2009; Beetz et al., 2018), ORO (Lemaignan et al., 2010), PMK (Diab et al., 2019) and CARESSES (Bruno et al., 2019). These works have explored and proven the relevance and usefulness of ontologies in robotics. However, they did not address the terminology defined in OCRA. Some other authors have focused on industrial robotic applications. Stenmark and Malec (2015), proposed the ROSETTA ontology, aimed at supporting reconfiguration and adaptation of robot-based manufacturing cells. Balakirsky (2015) implemented an ontology-based system for automatic recognition and adaptation to changes in manufacturing workflows. Stipancic et al. (2016), proposed to use a set of ontologies to semantically enrich the robots sensors data in order to enhance the decision making process in a multi-agent scenario. Chen et al. (2021), presented an ontology for automatic disassembly applications to represent terms related to

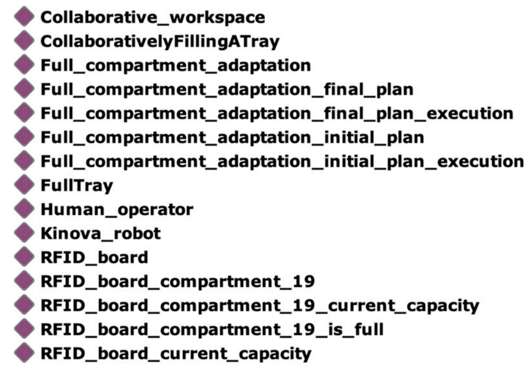

Fig. 7. Collaboratively filling a tray: ABox to answer competency questions about plan adaptation in Protege.

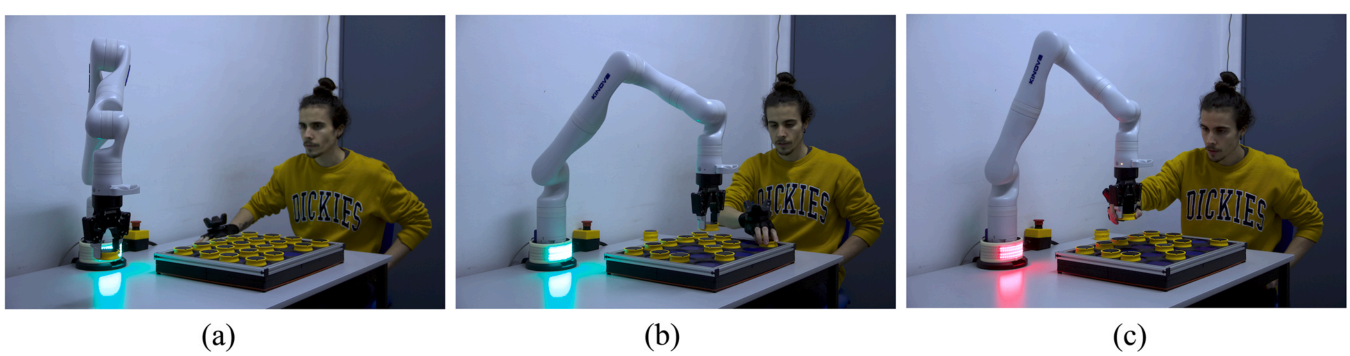

Fig. 5. Collaboration types. (a) Non-physical collaboration: the robot is selecting the compartment to place a token and the human monitors how the robot does its part of the plan. (b) Indirectly physical collaboration: the human and the robot move to place a token in different compartments without exchanging forces. (c) Directly physical collaboration: the human moves the robot exchanging forces while the robot remains in admittance mode.

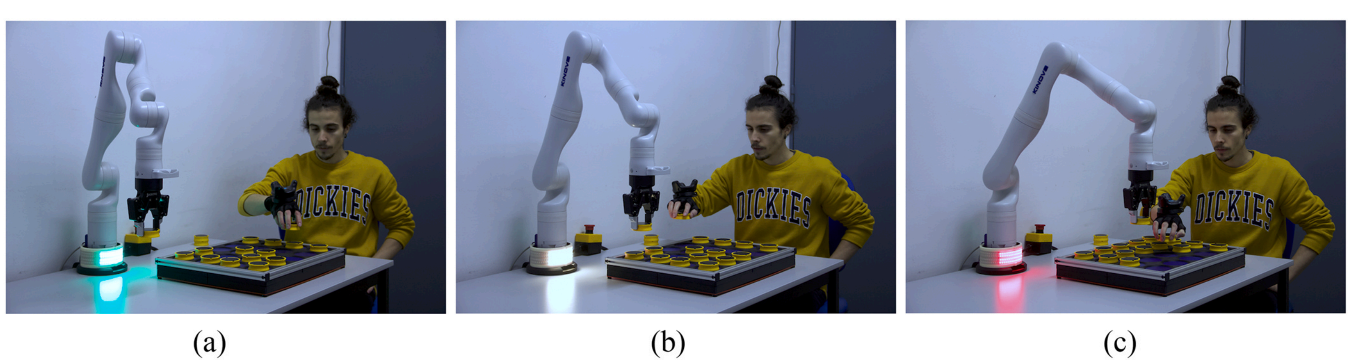

Fig. 6. Collaboration risks. (a) Low risk - green light: there is not any potential detected collision. (b) Medium risk - orange light: the robot has detected a possible collision. (c) High risk - red light: the detected collision is imminent.

(a) (b) (c)

Fig. 8. Plan adaptation to unforeseen events. In the shown sequence, (a) the human and the robot move towards the same compartment; (b) then the human fills the compartment; (c) the robot adapts its plan and moves to another free compartment.

Table 3
Analysis of OCRA's performance in a set of formalization's limit cases.

Case description	Classification	FOL formalization	OWL DL formalization
A human (h) and a robot share the plan and the goal during the plan's execution (e) but the human performs no activity.	Not a collaboration since only one of the agents (the robot) is active.	*dul.hasParticipant(e,h)* is violated, thus the case is excluded by the definition in FOL.	*dul.hasParticipant(e,h)* is violated, thus the case is excluded by the definition in OWL DL.
A human and a robot share the plan (p) and the goal during an event (e) in which they both perform activities but without executing the shared plan.	Not a collaboration since the event does not execute the shared plan.	*executesPlan(e,p)* is violated, thus the case is excluded by the definition in FOL.	*executesPlan(e,p)* is violated, thus the case is excluded by the definition in OWL DL.
A human and a robot have the same plan (p) and goal (g) during the time that both execute the plan, but the plan's goal is different from the shared goal.	Not a collaboration since the agents execute a plan to achieve a goal that is not shared.	*dul.hasComponent(p,g)* is violated, thus the case is excluded by the definition in FOL.	*dul.hasComponent(p,g)* is violated, thus the case is excluded by the definition in OWL DL.
A human and a robot share the plan during the time that its execution lasts (t). They also share the goal (g) but not during the whole execution, because the robot (r) changes its goal at some point.	Not a collaboration since the human and the robot do not share the goal during the whole execution of their plan.	*hasGoal(r,g,t)* is violated, thus the case is excluded by the definition in FOL.	*hasGoal(r,g)* holds some time, thus the case is not excluded by the definition in OWL DL. It might be solved by 'reification', introducing several relations *hasGoal_tp*(*r*, *g*), one for each instant (tp) in which the relation holds. This solution is activity dependent so we do not present it in the general definition.
An agent (robot) during the execution of plan (o) and due to a situation (s), realizes that there is a plan (n) that has the same goal and is better than the initial plan (o). However, the agent continues executing (f) the old plan (o).	Not a plan adaptation since the agent still executes the initial plan.	*executesPlan(f,n)* ∧ ¬*executesPlan (f,o)* are violated, thus the case is excluded by the definition in FOL.	*executesPlan(f,n)* ∧ ¬*executesPlan(f,o)* are violated, thus the case is excluded by the definition in OWL DL.
An agent (robot) during the execution of plan (o) decided to change and to execute another plan (n) which has the same goal. Nevertheless, the new plan (n) is not better than the original plan (o) due to the actual situation (s) realized after the execution of an initial part (i) of the original plan (o).	Not a plan adaptation since there is no situation that makes the new plan a better one.	*betterPlan(s,n,o)* is violated, thus the case is excluded by the definition in FOL.	*isBetterPlanThan(n,o)* is violated, thus the case is excluded by the definition in OWL DL.

processes, tools and production pieces such as fasteners. Although relevant for their domains, none of these works provided a formal definition for the concepts discussed in our work.

Of special interest is the work of Umbrico et al. (2020), who defined an ontology for human-robot collaboration. They focused on terminology which was mostly different to the notions defined in our work. Indeed, both ontologies could coexist and complement each other. The only overlap was regarding the notion of `Collaboration`. We think that our definition is stronger and more general, because it is based on a thorough analysis of how the concept was defined in the literature. Hence, it does not only represent our perspective but also a view shared by several works, including theirs. Furthermore, their ontology lacked other notions covered in OCRA such as `Collaboration Place`, or `Plan Adaptation`.

Finally, we can also find several works about ontologies for the industrial domain in general (Borgo et al., 2019; Karray et al., 2019; Liang, 2018; Liang, 2020; Mohd Ali et al., 2019; SampathKumar et al., 2019; Smith et al., 2019). Nonetheless, the content defined in OCRA cannot be found in any of them.

7. Conclusion

In this article we proposed OCRA, an Ontology for Collaborative Robotics and Adaptation. It has been built around two main notions: collaboration, and plan adaptation. The proposed definition of the main concepts is consistent with the state of the art. We also presented a formalization of the notions in FOL, to take advantage of its expressiveness, and provided a formalization in OWL DL for practical computational purposes. We have qualitatively validated the use of our ontology in a realistic case study in which a human and a robot share the execution of a task. First, we showed the capabilities of the ontology to answer a set of competency questions in a contextualized scenario. Second, we discussed how the formalization would work in some limit cases in which we purposely defined wrong instances of `Collaboration` and `Plan Adaptation`. Using OCRA, robots can formally represent and reason about their plan adaptations and collaborations in unstructured collaborative robotic scenarios. This work is a step forward to more reliable collaborative robots, and also to enhance the interoperability and reusability of the terminology in this domain. In the future, we aim to explore the use of the ontology in more tasks, in real industrial setups, and with more users. We also want to integrate OCRA in a system able to store a long-term knowledge base (or episodic memory). This memory could be used for robot behavior inspection in industrial scenarios, aiming for explainable collaborative robots. For instance, we could evaluate how adequate and understandable are the robot's explanations about their behavior. Furthermore, we could also use the episodic memory for learning tasks. For example, we might model the preferences of different users, or learn the structure of tasks to generalize to new ones.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been partially funded by MCIN/ AEI /10.13039/ 501100011033, Spain, under the project CHLOE-GRAPH (PID2020-119244GB-I00); by MCIN/ AEI /10.13039/501100011033, Spain, and by the "European Union NextGenerationEU/PRTR, Spain, under the project COHERENT (PCI2020-120718-2); and by the European Commission- NextGenerationEU, through CSIC's Thematic Platforms (PTI+ Neuro-Aging), Spain. S. Borgo is partially supported by the European project OntoCommons (GA 958371), Italy. A. Olivares-Alarcos is supported by the European Social Fund and the Ministry of Business and Knowledge of Catalonia, Spain, through the FI 2020 grant.

References

Ajoudani, A., Zanchettin, A.M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., Khatib, O., 2018. Progress and prospects of the human-robot collaboration. Auton. Robots 42 (5), 957–975.

Antoniou, G., van Harmelen, F., 2004. Web Ontology Language: OWL. In: Staab S., Studer R. (eds) Handbook on Ontologies, 1st. Springer, Berlin, Heidelberg, pp. 67–92.

Balakirsky, S., 2015. Ontology based action planning and verification for agile manufacturing. Robot. Comput. Integr. Manuf. 33, 21–28.

Barandiaran, X., Moreno, A., 2008. Adaptivity: from metabolism to behavior. Adapt. Behav. 16 (5), 325–344.

Barandiaran, X.E., Paolo, E.D., Rohde, M., 2009. Defining agency: individuality, normativity, asymmetry, and spatio-temporality in action. Adapt. Behav. 17 (5), 367–386.

Bauer, W., Bender, M., Braun, M., Rally, P., Scholtz, O., 2016. Lightweight robots in manual assembly-best to start simply. examining companies' initial experiences with lightweight robots., Tech. Frauenhofer-Institut für Arbeitswirtschaft und Organisation IAO, Stuttgart.

Bauer, A., Wollherr, D., Buss, M., 2008. Human-robot collaboration: a survey. International Journal of Humanoid Robotics 5 (01), 47–66.

Beetz, M., Beßler, D., Haidu, A., Pomarlan, M., Bozcuoğlu, A.K., Bartels, G., 2018. Knowrob 2.0 – a 2nd generation knowledge processing framework for cognition-enabled robotic agents, In: Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 512–519.

Booch, G., Rumbaugh, J., Jacobson, I., 2005. Unified Modeling Language User Guide. Addison-Wesley Object Technology Series, second ed. Addison-Wesley Professional.

Borgo, S., 2019. An ontological view of components and interactions in behaviorally adaptive systems. J. Integr. Des. Process. Sci. 23 (1), 17–35.

Borgo, S., Cesta, A., Orlandini, A., Umbrico, A., 2019. Knowledge-based adaptive agents for manufacturing domains. Eng. Comput. 35 (3), 755–779.

Borgo, S., Ferrario, R., Gangemi, A., Guarino, N., Masolo, C., Porello, D., Sanfilippo, E.M., Vieu, L., 2021. Dolce: a descriptive ontology for linguistic and cognitive engineering. Appl. Ontol. Prepr. 1–25 (preprint).

Borst, P., Akkermans, H., Top, J., 1997. Engineering ontologies. International Journal of Human-Computer Studies 46, 365–406.

Bruno, B., Chong, N.Y., Kamide, H., Kanoria, S., Lee, J., Lim, Y., Pandey, A.K., Papadopoulos, C., Papadopoulos, I., Pecora, F., 2019. The caresses eu-japan project: making assistive robots culturally competent. In: Italian Forum 2017 of Ambient Assisted Living. Springer, pp. 151–169.

Chacón, A., Angulo, C., Ponsa, P., 2020. Developing cognitive advisor agents for operators in industry 4.0. In: Martínez, L.R., Osornio, R.A., Prieto, M.D. (Eds.), New Trends in the Use of Artificial Intelligence for the Industry 4.0. IntechOpen, pp. 127–142.

Chen, W.H., Foo, G., Kara, S., Pagnucco, M., 2021. Automated generation and execution of disassembly actions. Robot. Comput. -Integr. Manuf. 68, 102056.

Diab, M., Akbari, A., Ud Din, M., Rosell, J., 2019. PMK-A knowledge processing framework for autonomous robotics perception and manipulation. Sensors 19 (5), 1166.

Dillenbourg, P., 1999. What do you mean by collaborative learning? In: Collaborative learning: cognitive and computational approaches. (Advances in learning and instruction series). Elsevier Science, Inc, pp. 1–19.

Fernández-López, M., Gómez-Pérez, A., Juristo, N., 1997. Methontology: from ontological art towards ontological engineering. In: Farquhar, A., Gruninger, M., (eds) AAAI 1997 Spring Symposium, Technical Report SS-97-06. AAAI Press, Menlo Park, California, pp. 34–40.

Fiorini, S.R., Bermejo-Alonso, J., Gonçalves, P., Pignaton de Freitas, E., Olivares Alarcos, A., Olszewska, J.I., Prestes, E., Schlenoff, C., Ragavan, S.V., Redfield, S., Spencer, B., Li, H., 2017. A suite of ontologies for robotics and automation [industrial activities]. IEEE Robot. Autom. Mag. 24 (1), 8–11.

Garcia-Camacho, I., Lippi, M., Welle, M.C., Yin, H., Antonova, R., Varava, A., Borras, J., Torras, C., Marino, A., Alenyà, G., Kragic, D., 2020. Benchmarking bimanual cloth manipulation. IEEE Robot. Autom. Lett. 5 (2), 1111–1118.

Gennari, J.H., Musen, M.A., Fergerson, R.W., Grosso, W.E., Crubézy, M., Eriksson, H., Noy, N.F., Tu, S.W., 2003. The evolution of protégé: an environment for knowledge-based systems development. Int. J. Hum. -Comput. Stud. 58 (1), 89–123.

Gervasi, R., Mastrogiacomo, L., Franceschini, F., 2020. A conceptual framework to evaluate human-robot collaboration. Int. J. Adv. Manuf. Technol. 108 (3), 841–865.

Gjorven, E., Eliassen, F., Aagedal, J. O., 2006. Quality of adaptation. In: Proceedings of the International Conference on Autonomic and Autonomous Systems (ICAS'06), pp. 9–9.

Glimm, B., Horrocks, I., Motik, B., Stoilos, G., Wang, Z., 2014. Hermit: an owl 2 reasoner. J. Autom. Reason. 53 (3), 245–269.

Gopinath, V., Johansen, K., Derelöv, M., Gustafsson, Åke, Axelsson, S., 2021. Safe collaborative assembly on a continuously moving line with large industrial robots. Robot. Comput. Integr. Manuf. 67, 102048.

Gruber, T., 1993. A translation approach to portable ontologies. Knowl. Acquis. 5 (2), 199–220.

Gualtieri, L., Rauch, E., Vidoni, R., 2021. Emerging research fields in safety and ergonomics in industrial collaborative robotics: a systematic literature review. Robot. Comput. -Integr. Manuf. 67, 101998.

Guarino, N., 1998. Formal ontology in information systems. In: Proceedings of the FOIS'98. IOS Press, Amsterdam, The Netherlands, pp. 3–15.

Guarino, N., Giaretta, P., 1995. Ontologies and knowledge bases: towards a terminological clarification. In: N, M. (Ed.), Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing (KBKSa95). IOS Press, Amsterdam, The Netherlands, pp. 25–32.

Guarino, N., Oberle, D., Staab, S., 2009. What is an ontology? In: Handbook on Ontologies, 2nd. Springer, Berlin, Heidelberg, pp. 1–17.

Hou, J., List, G.F., Guo, X., 2014. New algorithms for computing the time-to-collision in freeway traffic simulation models. Comput. Intell. Neurosci.

I. ISO, 2010. 12100 Safety of machinery - general principles for design - Risk assessment and risk reduction.

ISO, I. , 2011a. 10218 Robots and robotic devices-Safety requirements for industrial robots - Part 2: Robot systems and integration.

ISO, I. , 2011b. 10218 Robots and robotic devices-Safety requirements for industrial robots - Part 1: Robots.

Järvenpää, E., Lanz, M., Tuokko, R., 2016. Application of a capability-based adaptation methodology to a small-size production system. Int. J. Manuf. Technol. Manag. 30 (1/2), 67–86.

Karray, M.H., Ameri, F., Hodkiewicz, M., Louge, T., 2019. Romain: towards a bfo compliant reference ontology for industrial maintenance. Appl. Ontol. 14 (2), 155–177.

Kim, W., Peternel, L., Lorenzini, M., Babič, J., Ajoudani, A., 2021. A human-robot collaboration framework for improving ergonomics during dexterous operation of power tools. Robot. Comput. -Integr. Manuf. 68, 102084.

G.L. Kolfschoten, 2007. Theoretical foundations for collaboration engineering, Ph.D. thesis, Faculty of Technology Policy and Management. Delft University of Technology., Jaffalaan 5, 2628 BX Delft, the Netherlands, a (12).

Lemaignan, S., Ros, R., Mösenlechner, L., Alami, R., Beetz, M., 2010. Oro, a knowledge management platform for cognitive architectures in robotics. In: Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 3548–3553.

Lenat, D.B., Guha, R.V., 1990. Building Large Knowledge-Based System: Representation and Inference in the Cyc Project. Addison-Wesley, New York.

Levine, S.J., Williams, B.C., 2014. Concurrent plan recognition and execution for human-robot teams. In: Proceedings of the Twenty-Fourth International Conferenc on International Conference on Automated Planning and Scheduling, ICAPS'14. AAAI Press, pp. 490–498.

Levine, S.J., Williams, B.C., 2018. Watching and acting together: concurrent plan recognition and adaptation for human-robot teams. J. Artif. Intell. Res. 63, 281–359.

Liang, J.S., 2018. An ontology-oriented knowledge methodology for process planning in additive layer manufacturing. Robot. Comput. -Integr. Manuf. 53, 28–44.

Liang, J.S., 2020. A process-based automotive troubleshooting service and knowledge management system in collaborative environment. Robot. Comput. -Integr. Manuf. 61, 101836.
Lints, T., 2010. The essentials of defining adaptation. In: Proceedings of the 2010 IEEE International Systems Conference, pp. 113–116.
Liu, H., Wang, L., 2021. Collision-free human-robot collaboration based on context awareness. Robot. Comput. -Integr. Manuf. 67, 101997.
Martín H, J.A., de Lope, J., Maravall, D., 2009. Adaptation, anticipation and rationality in natural and artificial systems: computational paradigms mimicking nature. Nat. Comput. 8 (4), 757.
McBride, B., 2004. The Resource Description Framework (RDF) and its Vocabulary Description Language RDFS. In: Staab S., Studer R. (eds) Handbook on Ontologies, 1st. Springer, Berlin, Heidelberg, pp. 51–65.
Melchiorre, M., Scimmi, L.S., Mauro, S., Pastorelli, S.P., 2021. Vision-based control architecture for human-robot hand-over applications. Asian J. Control 23 (1), 105–117.
Mohd Ali, M., Rai, R., Otte, J.N., Smith, B., 2019. A product life cycle ontology for additive manufacturing. Comput. Ind. 105, 191–203.
Nikolaidis, S., Kwon, M., Forlizzi, J., Srinivasa, S., 2018. Planning with verbal communication for human-robot collaboration. ACM Trans. Human-Robot Int. 7 (3), 1–21.
Niles, I., Pease, A., 2001. Towards a standard upper ontology. In: Proceedings of the 2001 International Conference on Formal Ontology in Information Systems. ACM, pp. 2–9.
O. for Economic Co-operation, 2017. Development, PISA 2015 assessment and analytical framework: Science, reading, mathematic, financial literacy and collaborative problem solving, OECD Publishing.
O'Connor, M., Das, A., 2009. Sqwrl: A query language for owl. In: Proceedings of the 6th International Conference on OWL: Experiences and Directions - Volume 529, OWLED'09, CEUR-WS.org, Aachen, DEU, pp. 208–215.
Olivares-Alarcos, A., Foix, S., Alenyà, G., 2019. On inferring intentions in shared tasks for industrial collaborative robots. Electronics 8 (11), 1306.
Olivares-Alarcos, A., Beßler, D., Khamis, A., Goncalves, P., Habib, M.K., Bermejo-Alonso, J., Barreto, M., Diab, M., Rosell, J., Quintas, J., Olszewska, J., Nakawala, H., Pignaton, E., Gyrard, A., Borgo, S., Alenyà, G., Beetz, M., Li, H., 2019. A review and comparison of ontology-based approaches to robot autonomy. Knowl. Eng. Rev. 34, e29.
Oliveira, F.F., Antunes, J.C., Guizzardi, R.S., 2007. Towards a collaboration ontology. In: Proceedings of the 2nd Brazilian Workshop on Ontologies and Metamodels for. Software and Data Engineering.
Oxford-University, P., 2022. Compact Oxford English Dictionary of Current English. URL ⟨http://www.askoxford.com⟩.
Pan, M.K., Knoop, E., Bächer, M., Niemeyer, G., 2019. Fast handovers with a robot character: Small sensorimotor delays improve perceived qualities. In: Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 6735–6741.
Paolo, E.A. Di, 2005. Autopoiesis, adaptivity, teleology, agency. Phenomenol. Cogn. Sci. 4 (4), 429–452.
Pérez, J., Arenas, M., Gutierrez, C., 2006. Semantics and complexity of sparql. In: The Semantic Web - ISWC 2006. Springer, Berlin, Heidelberg, pp. 30–43.
A. Rector, N. Noy, Defining n-ary relations on the semantic web, W3C note, W3C, ⟨https://www.w3.org/TR/2006/NOTE-swbp-n-aryRelations-20060412/⟩ (Apr. 2006).
Rozo, L., Calinon, S., Caldwell, D., Jiménez, P., Torras, C., 2013. Learning collaborative impedance-based robot behaviors 27, 1422–1428.
SampathKumar, V.R., Khamis, A., Fiorini, S., Carbonera, J.L., Olivares-Alarcos, A., Habib, M., Goncalves, P., Li, H., Olszewska, J.I., 2019. Ontologies for industry 4.0. Knowl. Eng. Rev. 34 (e17).
Schlenoff, C., Prestes, E., Madhavan, R., Goncalves, P., Li, H., Balakirsky, S., Kramer, T., Migueláñez, E., 2012. An IEEE standard ontology for robotics and automation. In: Proceedings of the 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1337–1342.
Scimmi, L.S., Melchiorre, M., Troise, M., Mauro, S., Pastorelli, S., 2021. A practical and effective layout for a safe human-robot collaborative assembly task. Appl. Sci. 11 (4), 1763.
Silverman, B.G., 1992. Human-computer collaboration. Hum. -Comput. Interact. 7 (2), 165–196.
Smit, B., Wandel, J., 2006. Adaptation, adaptive capacity and vulnerability. Glob. Environ. Change 16 (3), 282–292.
Smit, B., Burton, I., Klein, R.J., Wandel, J., 2000. An anatomy of adaptation to climate change and variability. Clim. Change 45 (1), 223–251.
Smith, B., Ameri, F., Cheong, H., Kiritsis, D., Sormaz, D., Will, C., Otte, J.N., 2019. A first-order logic formalization of the industrial ontology foundry signature using basic formal ontology. In: The Joint Ontology Workshops, 10th International Workshop on Formal Ontologies meet Industry (FOMI). CEUR-WS.org.
Smullyan, R.M., 1968. First-Order Logic. Springer-Verlag Berlin Heidelberg.
Spyns, P., Tang, Y., Meersman, R., 2008. An ontology engineering methodology for dogma. Appl. Ontol. 3 (1–2), 13–39.
Stenmark, M., Malec, J., 2015. Knowledge-based instruction of manipulation tasks for industrial robotics. Robot. Comput. -Integr. Manuf. 33, 56–67.
Stipancic, T., Jerbic, B., Curkovic, P., 2016. A context-aware approach in realization of socially intelligent industrial robots. Robot. Comput. -Integr. Manuf. 37, 79–89.
Studer, R., Benjamins, V., Fensel, D., 1998. Knowledge engineering: Principles and methods. Data Knowl. Eng. 25 (1–2), 161–197.
Tenorth, M., Beetz, M., 2009. Knowrob - knowledge processing for autonomous personal robots. In: Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 4261–4266.
Terveen, L.G., 1995. Overview of human-computer collaboration. Knowl. -Based Syst. 8 (2–3), 67–81.
Umbrico, A., Cesta, A., Cortellessa, G., Orlandini, A., 2020. A holistic approach to behavior adaptation for socially assistive robots. Int. J. Soc. Robot. 12 (3), 617–637.
Umbrico, A., Orlandini, A., Cesta, A., 2020. An ontology for human-robot collaboration. Procedia CIRP 93, 1097–1102 53rd CIRP Conference on Manufacturing Systems 2020.
Vicentini, F., 2020. Terminology in safety of collaborative robotics. Robot. Comput. -Integr. Manuf. 63, 101921.
Villani, V., Sabattini, L., Loch, F., Vogel-Heuser, B., Fantuzzi, C., 2021. A general methodology for adapting industrial HMIs to human operators. IEEE Trans. Autom. Sci. Eng. 18 (1).
Zanchettin, A.M., Rocco, P., Chiappa, S., Rossi, R., 2019. Towards an optimal avoidance strategy for collaborative robots. Robot. Comput. Integr. Manuf. 59, 47–55.