ORKA: An Ontology for Robotic Knowledge Acquisition

Mark Adamik[1 [0000−0002−7977−3617]], Romana Pernisch[1,2 [0000−0001−8590−1817]], Ilaria Tiddi[1 [0000−0001−7116−9338]], and Stefan Schlobach[1 [0000−0002−3282−1597]]

[1] Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
[2] Discovery Lab, Elsevier, Amsterdam, The Netherlands

Abstract. Most autonomous agents operating in the real world use perception capabilities and reasoning mechanisms to acquire new knowledge of the environment, where perception capabilities include both the physical sensor devices and the software-based perception pipelines involved in the process. For autonomous agents to be able to adjust and reason over their own perception, knowledge of the sensors and the corresponding perception algorithms is required. We present the Ontology for Robotic Knowledge Acquisition (ORKA), that models the perception pipeline of a robotic agent by representing the sensory, algorithmic and measurement aspects of the perception process, thereby unifying the agent's sensing with the characteristics of the environment and facilitating the grounding process. The ontology is based on the alignment between SSN and OBOE, linked to external databases as additional knowledge sources for robotic agents, populated with instances from two different robotic use-cases, and evaluated using competency questions and comparisons to related ontologies. A proof of concept use-case is presented to highlight the potential of the ontology.

Keywords: Robotic Knowledge Acquisition · Robotic Perception · Ontologies

1 Introduction

Robots are increasingly used across various sectors due to advancements in AI and robotics. They range from service robots in restaurants and hospitals through home-use robots like vacuum cleaners to agricultural robot. Yet, their utility and effectiveness is hindered by a lack of common-sense knowledge and understanding of the world [12]. Presently, the common-sense knowledge robots use is implicitly embedded within specialised control programs designed for various robots and applications.

Fundamentally, the autonomy and behaviour of robotic systems are shaped by their ability to perceive their surroundings, as many of the decisions these agents make are based on the interpretation of data acquired through sensors. However, the data acquired through the sensors is essentially meaningless without the background knowledge that allows for the interpretation of this information. In order to efficiently organise the data and transform it into knowledge that could be acted on, the agent needs to have some prior information about what type of perception capabilities it is equipped with, which sensors expose them, and what aspects of the environment are observed by them.

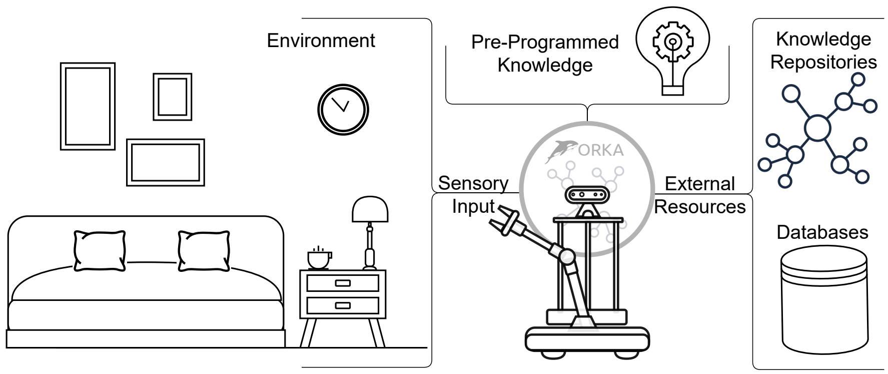

Fig. 1: The three main sources of knowledge for a robotic agent is the pre-programmed knowledge that the robotic system engineers equipped the robot with; the sensory information the robot acquires through the sensory data, and external sources the robot can connect to. The purpose of ORKA is to unify these knowledge sources.

Let us consider a robot equipped with two sensors: a 2D planar LiDAR for depth sensing and an RGB camera for vision. For a robot to interpret these signals, knowledge of the robotic engineer is needed to consider that the floating point values returned by the LiDAR represent distances represented (usually) in units of millimetres, and that the integer values of the camera represent pixel values corresponding to colour intensities on three channels. This knowledge, often implicit and only considered by system designers, is usually not directly available to the robot to reason over, thereby limiting its potential utility and its adaptability to unforeseen circumstances.

Similarly, considering the perception pipelines designed for these sensors, engineers are typically knowledgeable about the semantics of an algorithm's output – e.g. bounding boxes of generic, deep learning-based object recognition algorithms such as Yolo [27] indicate the location of the object, where the object class label represent the type of the object and the values represent the certainty of such classification. Outputs of a specifically-purposed gaze detector algorithm [3] instead indicate the detected face, as well as the yaw and pitch angles of the detected direction. To enable agents to autonomously process this information and independently select the optimal algorithms aligned with both the sensory abilities and the task at hand, awareness of their own perception capabilities is required.

Knowledge-enabled robotics [13] aims at supporting robotic agents to such type of reasoning using knowledge representation formalisms, to ultimately allow for a shared understanding between robotics agents and their environments. We distinguish three fundamental sources of knowledge for robotic agents as depicted in Figure 1: (i) *pre-programmed knowledge* that the robot is initially equipped with, and is designed by the engineers of the robot, also encompassing the knowledge representation and reasoning abilities such as ontologies or internal simulations [11,33] that robots use to infer new knowledge; (ii) *sensorial knowledge* gathered by the sensors and perception algorithms that create observations of the environment[3]; and (iii) the *external knowledge* that robots can access from databases [36] and knowledge graphs [14] to further improve their reasoning abilities. One of the most notable example of the latter source of knowledge is the Semantic Web, as one of the primary vision of it is to *"make the web more accessible to computers"*[6]. The problem of linking of robotic perception to the resources of the Semantic Web has been identified and formalised in our previous work [4].

While some studies have proposed models that represent parts of such knowledge in the form of ontologies (e.g. robot capabilities and sensory devices [19,24]), none of these attempts represent all the above described sources of knowledge, and unify the knowledge acquisition process of a robotic agent from sensors to understanding the environment. Furthermore, the represented characteristics of objects are somewhat limited in these ontologies [5]. With this in mind, we use the Ontology101 methodology [25] to define the Ontology for Robotic Knowledge Acquisition (ORKA). ORKA ties together the domains of sensors, perception processes, measurement properties of percepts and the properties of the perceived objects to jointly represent the domain of robotic perception. The contributions are two-fold. First, we present ORKA, a domain ontology for robotic knowledge acquisition, as well as the most important design choices made for it in the form of competency questions. Second, we model the robotic knowledge acquisition process from sensory data to real world entities, showing with two practical examples how such ontology-driven representation can benefit different use-cases. To the best of our knowledge, this is the first attempt to create an ontology that unifies and ties together the different sources of knowledge and their interrelations in artificial agents. The ontology and the related resources are available online[4].

2 Related Work

One of the most comprehensive knowledge-based robotics framework is KnowRob [33], which includes the Socio-physical Model of Activities (SOMA) [10], the Semantic Robot Description Language (SRDL) [19] and RoboSherlock [9,8]. Designed to aid robots performing manipulation activities in home environments, SOMA presents a very fine-grained model for objects and their social (e.g. cleaningness) and physical (e.g. colour) qualities. The SRDL module includes a detailed taxonomy of the different sensory capabilities, as well as some software categories. RoboSherlock is a cognitive vision system that provides an image processing-based perception pipeline to accommodate different perceptual processes. While KnowRob and its extensions provide an extensive and more detailed knowledge representing perception, they still lack a fine-grained taxonomy of characteristics entities can possess, the links of these characteristics to the robot sensory capabilities, and knowledge about the measurement standards of the sensory devices as well as the algorithms used to acquire these information.

Alternative knowledge-driven approaches include the Perception and Manipulation Knowledge (PMK) framework [24], which represents both object properties and

[3] We consider a "perception algorithm" any algorithmic process that results in new observations about characteristics of objects.

[4] https://github.com/Dorteel/orka

algorithms; the OpenRobots Ontology (ORO) [20], a knowledge management platform allowing cognitive robots to perform reasoning on previously acquired knowledge, and the Ontology-based Unified Robot Knowledge (OUR-K) [21], exhibiting integrating low-level sensory data with perceptual features (e.g. colour, texture, SIFT features), which are further associated with concepts and perception algorithms. The project seems however discontinued and no further details could be explored. There are some ontologies attempting to link sensory data to higher-level robot capabilities in specific domains, e.g. the culturally-aware assistive robots in the CARESSES project[5], including objects and qualities such as colour and size, but restricting the perception modelling to the audio (e.g. speed and pitch of the speech); the OROSU ontology [15] representing medical devices and their sensing measurements in the robotic surgery domain; and the ROSETTA ontology [32], where sensors and sensory characteristics are designed for industrial robotics tasks.

Other ontologies for sensing outside the robotics domain include the Semantic Sensor Network (SSN) [16] and its extension Sensor, Observation, Sample and Actuator (SOSA) ontology [17], both describing sensors and their observations, as well as sensing processes according to Semantic Web standards. While both ontologies offer a comprehensive overview of the sensory observation processes, they do not provide details about the specific sensors and corresponding algorithms employed by robots, as well regarding the characteristics of the environment the measurements are obtained about. The RDF Data Cube Vocabulary [2] allows to describe statistical data and their measurements, but places a greater emphasis on statistical data, and lacks the capabilities to describe sensory data accurately. Similarly, the Ontology of units of Measure and related concepts (OM) [29] provides a vocabulary to tie together different measurement units, with a focus on science and engineering [28]. The Extensible Observation Ontology (OBOE) [22] describes ecological observations of entities using measurement standards and characteristics. Its structure of the observations is easily transferable to the robotics domain, and the characteristics described in the ontology are quite comprehensive.

As shown, various ontologies incorporate aspects relevant to our problem (e.g. sensors, processes, characteristics), but none fully includes a suitable representation connecting the characteristics of an environment with the sensory devices used by a robot. Additionally, perception algorithms are hardly represented, and the algorithms' capabilities and the semantics of their outputs is missing. These ontologies also do not make use of external knowledge sources. To overcome these limitations, we present the Ontology for Robotic Knowledge Acquisition (ORKA) in the next section.

3 Ontology for Robotic Knowledge Acquisition

Following the Ontology101 [25] Ontology Engineering methodology, we start by defining the domain, purpose and intended use of the ontology articulated through a list of competency questions. Subsequently, we evaluate some of the ontologies of Section 2 for possible re-use. Finally, we present the structure of ORKA, as well as the Semantic Web Rule Language (SWRL) rules that augment its reasoning capabilities.

[5] http://caressesrobot.org/ontology/

3.1 Domain and Scope

The ontology serves as a shared vocabulary for researchers and engineers working in the domain of robotic perception and cognitive robotics, and defines the interrelations between the concepts used to describe sensors, subjects of sensors and the algorithms that operate on the sensory data. ORKA also includes an integration of additional knowledge sources, that could be used to improve the perception of the autonomous agent with common-sense information (as shown later in Section 5). The domain represented by ORKA is the complete process of knowledge acquisition for autonomous robotic agents, with a special focus on the link between robotic perception and environment.

We show the primary components the ontology should encompass using a few practical examples. An ideal ontology would allow autonomous robotic agents to consider that the sensors provide information about specific aspects of the environment, and the entities contained within. For example, some sensors such as IMU-s or encoders provide information about the state of the robot, and not of the objects it perceives; cameras provide information about objects that are detected, but not their measurements; LiDARs provide information about distances, but not colours. At the level of algorithms, not all sensory data and processes need to be considered at all times. For example, gaze-detection algorithms would only be required in contexts where a human is present, and sensors used for mapping are usually not required during manipulation tasks. Defining contexts for the different sensory applications could therefore help the agent decide which of the sensory data is relevant for the given task. Additionally, algorithms come in different model sizes, and ideally an agent should have the knowledge to deploy the most suitable model depending on the task at hand. Finally, in order to give a meaning to the sensor readings in the form of entities of the real-world, it is important that the measurement systems are also specified within the ontology. These entities can be also aligned with external, common-sense knowledge sources to acquire further knowledge: e.g. an object recognition algorithm returning contradicting or incorrect information could use a knowledge source such as WikiData [35] and DBpedia [7] to augment the object detection algorithm with the necessary information.

Restricting scope, we limit our investigation to the perception of physical objects, and refer the inclusion of other entities, such as events and actions, to future work.

Competency Questions. Using the above defined use-cases and scope limitations, we derive a set of competency questions to be answered by ORKA.

CQ1: Given a robotic agent, what sensors and perception algorithms is it equipped with?
CQ2: Which characteristics of the environment do the sensors and their associated algorithms from CQ1 observe?
CQ3: What units of measurement is the data from the sensors and perception algorithms in CQ1 provided in?
CQ4: What observable and observed characteristics do given entities possess?
CQ5: What characteristics do the sensors from CQ1 possess?
CQ6: What characteristics do the algorithms from CQ1 possess?
CQ7: Which algorithm is the most suitable for a given context?

CQ8: What external knowledge sources are available to the robot and what do they describe?
CQ9: What characteristics of given entities are described in the external knowledge sources?

CQ1-3 define the foundational concepts the ontology aims to connect, namely the sensors, algorithms, the characteristics, and the measurement standards. **CQ4** focuses on the characteristics which are implied by the existence of entities. **CQ5** and **CQ6** aim to describe the characteristics of sensors and algorithms respectively. Derived from the use-cases described earlier, these competency questions focus on the characteristics that are relevant to the knowledge acquisition process. This is reflected by **CQ7**, which requires the characteristics of **CQ5** and **CQ6** to determine the utility of the sensors and algorithms in a given context. As the comprehensive representation of a context is outside of the current scope of ORKA, the simplified representation of context includes (i) a task at hand, (ii) the perceived objects, and (iii) the objects required for completing the task. We disregard the particular task and instead use the required objects as a substitute for it. **CQ8** and **CQ9** address the third source of knowledge described in Section 1, i.e. the inclusion of external sources of knowledge in the knowledge-acquisition process. **CQ8** aims to describe the knowledge sources available to the robot, whereas **CQ9** focuses on the potential knowledge that these sources could offer.

3.2 Re-using existing ontologies

With the domain and scope defined, we consider the re-use of existing ontologies. As mentioned in Section 2, existing ontologies allow the modelling of sensors [16,17,19], observations [16,17,22] or tasks and activities [33], but none is comprehensive and flexible enough represent our robotic knowledge acquisition process sufficiently.

The sources that describe most of the required components are the SSN [16] ontology when considered together with the system capabilities module of SSN, which aims at extending SSN to capture the capabilities of sensors and measurement. The choice for reusing SSN was further reinforced with the fact that an alignment module is provided for OBOE [22], which in turn describes the characteristics the observations concern, and the associated units of measurements. Although some components are not included in the alignment of SSN and OBOE (i.e. the specific sensors, the coupling between the sensors and the measured characteristics, the specific algorithms that comprise the perception pipelines, the external knowledge sources, etc.), it is deemed as the best candidate to serve as a backbone of ORKA.

3.3 Core structure

We followed a top-down approach, starting with the definitions of the most general concepts. These considerations resulted in a core ontology (Figure 2), specified below.

Robots & Sensors. In accordance with SSN, in ORKA robots are considered a special class of `sosa:Platform`, that host `sosa:Sensor` devices and have specific `sosa:Procedures` implemented on. To address **CQ1**, sensors are also defined as a subclass of the `oboe:Entity`

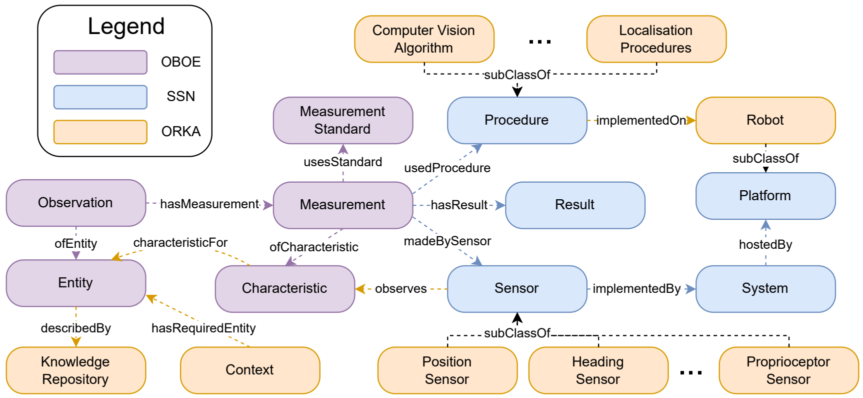

Fig. 2: A simplified view of the core classes of the ORKA ontology as well as the OBOE and SSN alignment.

class, and sensor instances are linked to an instance of the `orka:Robot` class using the `sosa:hosts` predicate. SSN is extended with a detailed classification of sensors following [30]. A sensor can be either a `ProprioceptorSensor`, i.e. a device that can measure different characteristics related to the body of the agent that equips the sensors, or an `ExteroceptorSensor`, which measures the external environment. The former includes encoders that measure the position of the motor devices, or sensors that measure the orientation or acceleration of the robot, such as tilt sensors or IMU devices. Using this classification the robot can infer that the sensory information concerns robots:

$$\text{ProprioceptorSensor} \sqsubseteq \forall \text{observesCharacteristic}.(\forall \text{characteristicFor}.\text{Robot})$$

An additional classification is provided based on the characteristic of the environment or the objects that they measure (e.g. position, heading, etc). As an example, the sensory activation of a bumper sensor always implies the `Location` characteristic of an entity.

$$\text{Bumper} \sqsubseteq \forall \text{observesCharacteristic}.\text{Location}$$

In determining the proprioception abilities of sensors, some adjustments are made. For example, some sensors indicated as exterioceptors in [30], such as compass and encoders, are modelled as proprioceptors in ORKA, as they indicate the orientation or position of the robot (or components of the robot).

To be able to address **CQ5**, the most important characteristics that influence the knowledge acquisition capabilities of the sensors are inserted as specific data properties. Some of these properties, such as `range`, `resolution`, `accuracy`, `precision` are included in the system capabilities module of SSN, while others, such as `sampling rate` of the sensors needed to be added. These properties are appended directly to the bottom levels of the `Sensor` class, where the specific models of sensors are represented, as general class axioms. Simpler sensors such as a thermometer or a bumper provide the characteristics of the entity without the necessity of a specialised perception algorithm to be utilised. Sensors that provide more complex information about the environment, such as a camera, require a specialised (often computer vision) algorithm, here collectively called `Procedures`, that allows for the observation of additional characteristics that the sensor does not provide on its own.

Procedures. Procedures in perception pipelines, acting on sensor data, significantly affect autonomous agents' capabilities. To capture this impact, we incorporate both algorithms and deep learning models together with their properties into the ontology. In order to address **CQ6**, certain characteristics (i.e. model size, inference speed, memory usage) and model versions of existing algorithms are included. For `ObjectDetection` algorithms, we currently use the official values of YoloV5[6] as implemented on our LoCoBot mobile robot, and represent these as data property assertions on the class level, in the form of general class axioms. These characteristics also play a role in answering **CQ7**. Depending on the context, models with higher detection speed could be preferred in some scenarios, whereas higher precision could be required in others. In the current version of the ontology, a greater emphasis is placed on the `ComputerVisionAlgorithm` class, which follows an application-based taxonomy [18]. Future work will be aimed at expanding it to a wider variety of algorithms.

Entity. The entities correspond to the phenomena being observed. In the case of robotic agents situated in a real-world environment, the properties of these entities could help distinguish between individuals, and address cognitive robotics problems such as object permanence and occlusion. Following OBOE, in ORKA entities are the main subjects of the observations, holding observable and inherent characteristics that describe the objects (**CQ4**) through the `hasCharacteristic` predicate. In ORKA, instances of `Sensor` and `Robot` are also considered as entities.

Observations & Measurements. Observations are made of entities belonging to the `Entity` class. In case of a sensor, an `Measurement` is produced that serves as an input to a `Procedure`. In the case of an object detection algorithm, the observation concerns a single entity, which has multiple corresponding `Measurements` involving the label and bounding-box that the procedure provides. Measurements concern a single `Characteristic`. A precision (where applicable) can also be assigned, as it can be inferred from the *precision* characteristics of the sensor device performing the measurement. The measurements also correspond to a measurement standard.

Measurement Standards. The sensory knowledge acquisition process starts with the recording of the observations of the physical world. In order to address **CQ3**, the measurement units contained within the ontology are adopted from OBOE. However, as OBOE does not include some units that are relevant for robotic perception (e.g. pixel values or binary events such as a switch or bumper), ORKA also incorporates these as measurement standards. Finally, instead of organising units into *base units* and *derived units* as in [22,29], ORKA defines a hierarchy of unit classes based on the characteristics they define.

[6] https://pytorch.org/hub/ultralytics_yolov5/

The `oboe:usesStandard` data property is used to link measurement standards to their corresponding measurements. In order to derive the measurement standards associated with the different sensors and procedures, the current implementation of ORKA uses SWRL rules [1]. The reason behind this modelling choice is that sensors are difficult to describe in general terms regarding their outputs, as different implementations of the same sensor type might produce outputs corresponding to different measurements. For example, a specific model of ultrasonic sensor might be implemented so that its output is the raw data in terms of time taken for the sound to reflect, while another model might also implement the conversion from the time taken to estimate the distance. Furthermore, some procedures, such as object recognition algorithms, produce multiple measurements corresponding to different units of measure, and therefore a single unit cannot be associated with a single procedure.

Characteristics. Given **CQ2**, a core aspect of ORKA is to represent characteristics of an object, that could be considered common-sense (i.e. where an object was observed, when it was observed or what it looks like). Given the lack of a comprehensive catalogue of the specific qualities (such as size, colour, shape, weight, etc.) that a particular endurant might have [23], we use three main sources of characteristics to include in ORKA. Firstly, we adopt the Characteristics defined in OBOE. Secondly, we examine the set of current perception algorithms, and sensors as well as their measurements to identify some of the main object characteristic to include in ORKA. This includes `ObjectType` (corresponding to the label), and `Location` as the output of the YOLO algorithm, or `Size` and `Colour` as the output of the point cloud processing algorithm that processes the images produced by the depth camera. This category also includes the measurement units discussed above, where we consider the SI quantity dimensions described in [34] (e.g. length, mass, time), and their derived characteristics (e.g. height, width, depth, weight, age, etc.). Lastly, we include some of the properties that external knowledge sources utilise to characterise objects, and which could be acquired from an external knowledge base. Specifically, WikiData [35] was used to examine additional characteristics, such as material, density, hardness and names of different shades of colours.

Observable characteristics are divided into two main categories: those related to objects and those related to the environment, influencing object perception (e.g. brightness impacting colour). The categories include the `SpatialCharacteristic` class, which outlines an object's location and orientation within a specific reference frame, and the `VisualCharacteristic` class, which defines an object's colour and pattern. This division suggests that robots without visual sensors should not be assigned tasks requiring the observation of visual characteristics. For example, a TurtleBot with only LiDAR is not suited for tasks involving visual characteristics, and its tasks should be adjusted or performed by robots equipped with the necessary sensors.

External Knowledge Sources. To be able to answer **CQ8** and **CQ9**, external knowledge sources, and links to these sources should be described. Currently, we limit our investigation to common-sense databases part of the Semantic Web that could be useful for robotic agents to operate in an environment, and use WikiData and DBpedia to serve as a proof-of-concept external knowledge graphs. As these sources contain information about general characteristics of the class of the entities being observed (e.g. colour, material, weight, etc.), but not about specific instances of entities that exist around the robotic agent, links are established between the characteristics included in ORKA and their counterpart in the external knowledge source. We link the resources using the `hasDBpediaURI` and `hasWikiDataURI` data properties. Furthermore, the `sparqlEndpoint` data property is used to provide some information on how to query these external knowledge sources. In the current version of ORKA, the corresponding links have been established manually. Automatic entity linking procedures could be considered in future work.

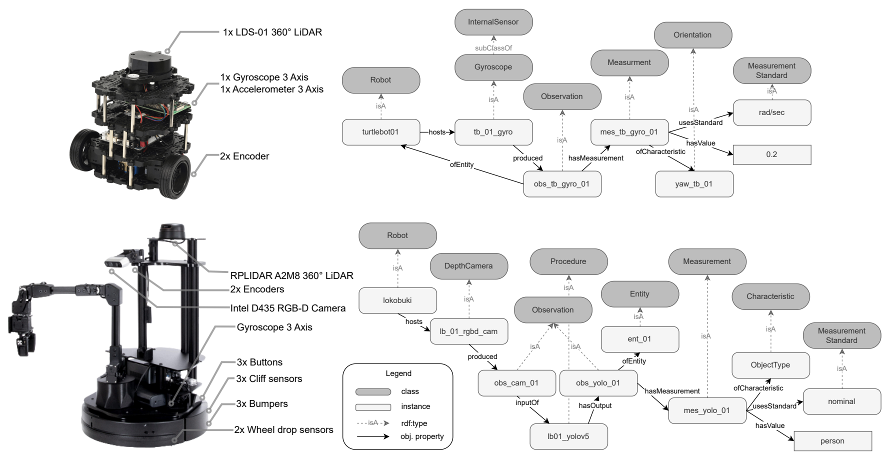

Fig. 3: A depiction of a TurtleBot 3 Burger[1] (top) and a LocoBot WX250s[2] (bottom) and their sensors, together with instances of sensory readings represented in the ontology. As the gyroscope of the TurtleBot is a proprioceptor, a SWRL rule automatically derives that the observations correspond to the entity that hosts the sensor. For the purposes of this paper, the manipulator of LocoBot is excluded from the consideration.

Individuals. The last step is the population of the ontology with individuals. We consider two mobile robots and their perception abilities to be instantiated: a LocoBot WX250s, and a TurtleBot3 Burger. The two robots are equipped with a total of 21 sensors and four procedures, a YoloV5 object detector, a gaze detector based on L2CS-Net [3], a point cloud segmentation algorithm and a Simultaneous Localisation and Mapping (SLAM) algorithm. A depiction of the robots as well as the corresponding sensors and an example of the sensory process can be seen in Figure 3. For every sensor

[1] https://emanual.robotis.com/docs/en/platform/turtlebot3/features/
[2] https://www.trossenrobotics.com/locobot-wx250.aspx

and algorithm within the ontology, instances of test observations have been added to demonstrate and validate the competency questions.

3.4 SWRL Rules

To address **CQ4** and **CQ7**, ORKA needs to infer certain relationships, such as which algorithms can detect the `ObjectType` and other characteristics of the given entities (**CQ4**) or which entities are required for a context (**CQ7**). As OWL-DL does not allow for the desired expressivity, we use SWRL rules [1] to infer the required information. The rules can be divided into rules concerning sensory data, and rules concerning procedures.As an example for the former, if the sensor is a `ContactSensor` and produces an `Observation` with a `Measurement ofCharacteristic Location`, then it should also be inferred that, since the measurement range of any `ContactSensor` is 0, the value of the location should be equated with the given sensor's `Location`. An example for the latter is the rule concerning **CQ7**, which aids in the deduction of the `canDetect` object property between a `Procedure` and an `Entity`. This involves establishing a connection between an entity required for a task, instantiated as an instance, and potential algorithms capable of completing the task by detecting the specified entity. The complete list of implemented SWRL rules with comments is available in the online repository[7].

4 Evaluation

The current version of ORKA contains 443 classes, 74 object properties and 39 data properties. The ontology is populated with 54 instances, representing the two robots, their capabilities, observations and measurements, as well as some of the observed entities. The competency questions that articulate the requirements of the ontology are evaluated using several SPARQL queries and the Pellet reasoner [31] . The queries were evaluated using the individual instances of the two robotic agents presented in Section 3. A complete list of the queries as well as the results can be found in the repository.

Listing 1 illustrates the query for **CQ1** which returns the sensors a robotic agent is equipped with, and the implemented perception algorithms. The query is used to give a complete list of the sensory knowledge acquisition capabilities of any robotic agent.

```
SELECT ?s ?r ?a
WHERE { ?r sosa:hosts ?s .
        ?a orka:implementedOn ?r}
```

Listing 1: Query corresponding to CQ1.

In practice, the list could be used by other robotic agents in a multi-agent setting, or by a user to assess which robot might be appropriate in which situation. The queries

[7] https://github.com/Dorteel/orka

corresponding to **CQ3** and **CQ8** follow a similar pattern, with entities being returned that have a `hasMeasurementStandard` and `describedBy` properties defined respectively. Listing 2 presents a more complex query corresponding to **CQ2**, and returns the individuals of the class `Characteristic` corresponding to sensors and procedures answered to **CQ1**.

```
SELECT ?item ?charType
WHERE {     {?a orka:implementedOn ?r .
             ?a orka:observesCharacteristic ?c .
             ?c a ?charType .
             BIND(?a AS ?item)}
    UNION   {?r sosa:hosts ?s .
             ?s orka:observesCharacteristic ?c .
             ?c a ?charType .
             BIND(?s AS ?item) }}
```

Listing 2: Query for the characteristics measured by sensors and algorithms (CQ2).

When evaluating competency questions **CQ4-6**, the sub-properties of `hasObservableCharacteristic`, `hasSensorCharacteristic` and `hasAlgorithmCharacteristic` are queried respectively. In a practical scenario, specific characteristics such as `VisualCharacteristics` of entities, `maxRange` of a sensor or `DetectionSpeed` of a procedure would be the subject of the query.

To determine the context in which certain perception algorithms are better than others, as per **CQ7**, two simple contexts are presented. A *HRI-Dialogue* context has a human face as a required entity, whereas a *Fetch-Object* context has the object to be fetched as the required entity. The query of Listing 3 formulates this question and returns the entities that match the description of the context, except where none of the algorithms can detect the required entity. Context requirements can be modified to select among algorithm characteristics, such as opting for an object detector version that prioritises higher detection speed over accuracy.

```
SELECT ?c ?e (COALESCE(?a, "None") AS ?alg)
WHERE { ?c orka:hasRequiredEntity ?e .
        OPTIONAL {
            ?a orka:implementedOn ?r .
            ?a orka:canDetect ?e .}}
```

Listing 3: Query used to evaluate CQ7.

Table 1: Ontology comparison with respect to CQ1 – CQ9. As KnowRob represents an infrastructure consisting of multiple ontologies, the aggregated capabilities are considered. Checkmarks in parentheses indicate partial answering of CQ-s.

Ontology	CQ1	CQ2	CQ3	CQ4	CQ5	CQ6	CQ7	CQ8	CQ9
KnowRob [33]	✓	✓		✓	✓	✓	✓		
PMK [24]	✓			✓	✓				
CARESSES				✓					
ROSETTA [32]	✓	✓	✓	✓	✓	✓			
ORO [20]				✓					
OROSU [26]	✓								
SSN [16]	✓	✓			✓	✓			
ORKA	✓	✓	✓	✓	✓	✓	✓	✓	(✓)

CQ9 necessitates listing the characteristics of entities from an external knowledge source, useful for correcting mislabelled entities by comparing the detected entity's characteristics with those described in the external source. The ontology is not able to evaluate this competency question by relying only on SPARQL queries, as it would require a dynamic construction of URI's within the query. However, using the `describedBy` property as in **CQ8** allows for the access of these external resources using an additional Python script that incorporates SPARQL queries.

Comparison. A comparison was performed to evaluate ORKA with respect to the other publicly available ontologies. Table 1 shows the coverage of the CQs not just of ORKA (last row), but also of some of the related ontologies described in Section 2. ORKA answers all but the last CQ completely, making it the most fitting ontology for the envisioned domain and scope. KnowRob [33] and ROSETTA [32] is able to answer six CQs, making them the next best-performing ontologies. PMK [24] and SSN/SOSA [17] can each answer four CQs, and ORO [20], OROSU [15] and CARESSES are only able to address one CQ. With the low coverage of CQs by these domain-specific ontologies, creating alignments with ORKA to extend the knowledge acquisition capabilities can be envisioned, which could ultimately serve as a unifying vocabulary for robot capabilities.

Finally, we compare ORKA and the relevant domain ontologies in terms of size. Table 2 includes the number of classes and instances for the sensors, algorithms, characteristics, measurement units, as well as sensor and algorithm characteristics that were considered within each ontology. These values are acquired using the DL Query plugin of the Protégé ontology development software.

Recall that ORKA is in its preliminary version, and the number of its individuals is expected to increase. We consider the ones we have enough to prove the ontology generalisability to a further extent. Further robot scenarios outside of the ones described in this paper are left for future work.

Threats to Validity. The competency questions that guided the development process were driven by the use-cases, and do not yet cover the entirety of the robotic knowledge acquisition domain. Currently, ORKA is limited to the individuals representing

Table 2: Number of classes and instances (in parentheses) of ORKA and comparison ontologies.

	KnowRob [33]	PMK [24]	CARESSES	ROSETTA [32]	ORO [20]	OROSU [15]	ORKA (ours)
Sensors	21	4 (4)	-	22 (69)	-	4 (4)	37 (18)
Algorithms	63	14 (7)	-	45 (61)	-	13 (5)	22 (4)
Characteristics	27	47 (40)	1	78 (245)	3 (13)	-	120 (11)
Measurement Units	-	-	-	2 (2)	-	-	17
Sensor Characteristics	6	2 (2)	-	13 (86)	-	-	9
Algorithm Characteristics	1	-	-	-	-	-	8
External Sources	-	-	-	-	-	-	2

the two robots and their associated sensory devices. For the ontology to be more complete, more links between characteristics and the sensor devices outside of the individuals in the ontology should be included. Moreover, the comparison only considers the freely available versions of the considered ontologies. More up-to-date versions could appear that would render the comparison out-of-date. ORKA's interoperability could be improved by introducing alignment modules with ontologies within and outside the domain. Lastly, our design methodology does not include evaluation with external validators, which also poses a threat to the validation of ORKA. Yet, we show how ORKA can support the well-known issue of robotics knowledge-driven perception, enabling further studies and first prototypes in the future.

5 Example Use-Case

This section provides a proof-of-concept with ORKA, where a robot detects different objects, and using an external knowledge graph (i.e. WikiData) verifies and corrects the label provided by the object detection algorithm, performing the perceived-entity linking task [4].

The example concerns the task of fetching an orange for the user from a basket. With the use of a camera, an object detection algorithm and a colour detection procedure, an observation graph is produced containing the entities and their recognised properties. As the fetching task contains an object that is recognisable with several of the available object detection algorithms, the one with the fastest inference speed is chosen. However, as the object detection algorithm incorrectly recognises two of the fruits (a grapefruit and a lemon) as an orange, a colour detection procedure is initiated to further distinguish the objects. The colour detection algorithm utilises a modified MASK-RCNN segmentation algorithm to calculate the mean pixel value for each segment, and consults ORKA to retrieve the corresponding closest available colour.

The robot associates each entity with a colour, and uses the corresponding colour property (`wdt:P462`) of WikiData to refer to the colour of an orange (accessed through the `hasWikiDataURI` data property) and therefore disambiguate which entity has the colour of an orange as described in WikiData. An overview of the process as well as the outputs are shown in Figure 4. This example serves as a proof-of-concept, showcasing

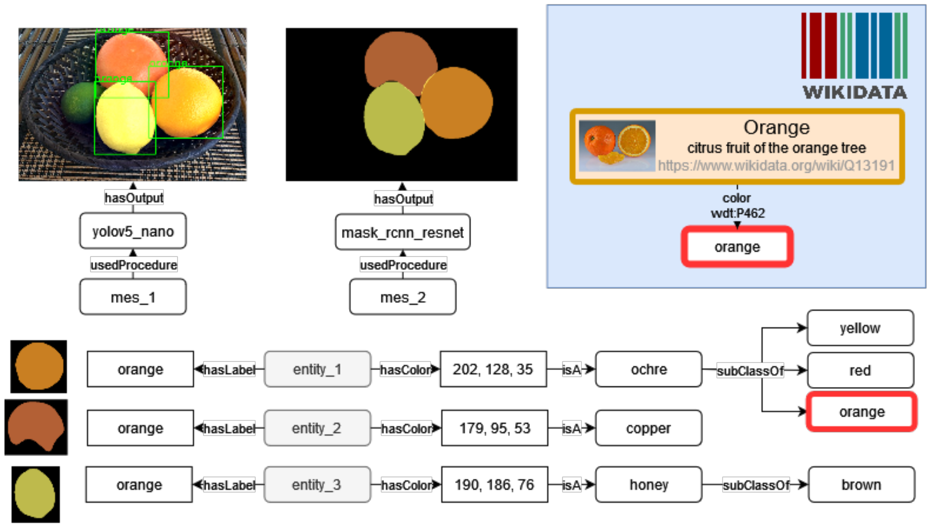

Fig. 4: A simplified depiction of the robot view, the created observation graphs concerning the entities, and the matched property. As only one entity has the colour matching the description found in the external knowledge source, the entity is chosen.

the utility and potential of ORKA to represent the knowledge and reasoning capabilities of the agents.

6 Conclusion and Future Work

We have presented the Ontology for Robotic Knowledge Acquisition (ORKA), a model that represents the sensory and algorithmic capabilities of robotic agents with regard to the perception of the environment. We have shown that the model is capable of capturing several characteristics of objects, and allows for the linking of object classes and characteristics to external knowledge sources. The reasoning of the model has been reinforced with SWRL rules to allow for the automatic inference of characteristics captured by the algorithms. The model is evaluated using the competency questions formulated as SPARQL queries, and a proof of concept is presented that showcased the potential of ORKA. In the future, we intend to extend the model to address problems such as the representation of time, using e.g. multiple descriptions of the same location. In order to promote re-usability and compatibility, alignment modules will be provided to allow other knowledge-driven robotic systems to use and integrate ORKA. A modularisation of ORKA will be introduced, which allows users and developers to focus on the aspects relevant to their respective applications. Furthermore, we plan to use ORKA in combination with physical robots, and test the ontology reasoning capabilities as more observations and measurements are made. We encourage experts and users in the field of knowledge representation and robotics to use and revise our model so that ORKA can become a community effort in the future.

References

1. SWRL: A Semantic Web Rule Language Combining OWL and RuleML. https://www.w3.org/submissions/SWRL/
2. The RDF Data Cube Vocabulary. https://www.w3.org/TR/vocab-data-cube/ (2014)
3. Abdelrahman, A.A., Hempel, T., Khalifa, A., Al-Hamadi, A., Dinges, L.: L2CS-Net : Fine-Grained Gaze Estimation in Unconstrained Environments. In: 2023 8th International Conference on Frontiers of Signal Processing (ICFSP). pp. 98–102 (Oct 2023). https://doi.org/10.1109/ICFSP59764.2023.10372944
4. Adamik, M., Pernisch, R., Tiddi, I., Schlobach, S.: Advancing robotic perception with perceived-entity linking. In: The Semantic Web - ISWC 2024 - 23rd International Semantic Web Conference, Baltimore, The United States of America, November 11-15, 2024, Proceedings. Lecture Notes in Computer Science, Springer (2024)
5. Adamik, M., Schlobach, S.: Towards a definition and conceptualisation of the perceived-entity linking problem. In: RobOntics 2023 Ontologies for Autonomous Robotics 2023. CEUR Workshop Proceedings, CEUR-WS (2023)
6. Antoniou, G., Groth, P., van Harmelen, F., Hoekstra, R.: A Semantic Web Primer, 3rd Edition. MIT Press (2012)
7. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: DBpedia: A Nucleus for a Web of Open Data. In: Aberer, K., Choi, K.S., Noy, N., Allemang, D., Lee, K.I., Nixon, L., Golbeck, J., Mika, P., Maynard, D., Mizoguchi, R., Schreiber, G., Cudré-Mauroux, P. (eds.) The Semantic Web. pp. 722–735. Lecture Notes in Computer Science, Springer, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76298-0_52
8. B'alint-Bencz'edi, F., Worch, J.H., Nyga, D., Blodow, N., Mania, P., M'arton, Z., Beetz, M.: RoboSherlock: Cognition-enabled Robot Perception for Everyday Manipulation Tasks. ArXiv (Nov 2019)
9. Beetz, M., Balint-Benczedi, F., Blodow, N., Nyga, D., Wiedemeyer, T., Marton, Z.C.: RoboSherlock: Unstructured information processing for robot perception. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 1549–1556. IEEE, Seattle, WA, USA (May 2015). https://doi.org/10.1109/ICRA.2015.7139395
10. Beßler, D., Porzel, R., Pomarlan, M., Vyas, A., Höffner, S., Beetz, M., Malaka, R., Bateman, J.: Foundations of the Socio-physical Model of Activities (SOMA) for Autonomous Robotic Agents (Nov 2020). https://doi.org/10.48550/arXiv.2011.11972
11. Blum, C., Winfield, A.F.T., Hafner, V.V.: Simulation-based internal models for safer robots. Frontiers Robotics AI **4**, 74 (2017). https://doi.org/10.3389/FROBT.2017.00074, https://doi.org/10.3389/frobt.2017.00074
12. Davis, E., Marcus, G.: Commonsense reasoning and commonsense knowledge in artificial intelligence. Communications of the ACM **58**(9), 92–103 (Aug 2015). https://doi.org/10.1145/2701413
13. Engel, U. (ed.): Robots in Care and Everyday Life: Future, Ethics, Social Acceptance. SpringerBriefs in Sociology, Springer International Publishing, Cham (2023). https://doi.org/10.1007/978-3-031-11447-2
14. Fischer, L., Hasler, S., Deigmoller, J., Schnurer, T., Redert, M., Pluntke, U., Nagel, K., Senzel, C., Ploennigs, J., Richter, A., Eggert, J.: Which tool to use? Grounded reasoning in everyday environments with assistant robots. Proceedings of the 11th International Workshop on Cognitive Robotics (2018)
15. Gonçalves, P.J., Torres, P.M.: Knowledge representation applied to robotic orthopedic surgery. Robotics and Computer-Integrated Manufacturing **33**, 90–99 (2015). https://doi.org/https://doi.org/10.1016/j.rcim.2014.08.014, citation Key: orosu

16. Haller, A., Janowicz, K., Cox, S.J., Lefrançois, M., Taylor, K., Le Phuoc, D., Lieberman, J., García-Castro, R., Atkinson, R., Stadler, C.: The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation. Semantic Web **10**(1), 9–32 (Dec 2018). https://doi.org/10.3233/SW-180320
17. Janowicz, K., Haller, A., Cox, S.J.D., Phuoc, D.L., Lefrancois, M.: SOSA: A Lightweight Ontology for Sensors, Observations, Samples, and Actuators. Journal of Web Semantics **56**, 1–10 (May 2019). https://doi.org/10.1016/j.websem.2018.06.003
18. Khanday, N.Y., Sofi, S.A.: Taxonomy, state-of-the-art, challenges and applications of visual understanding: A review. Computer Science Review **40**, 100374 (May 2021). https://doi.org/10.1016/j.cosrev.2021.100374
19. Kunze, L., Roehm, T., Beetz, M.: Towards semantic robot description languages. In: 2011 IEEE International Conference on Robotics and Automation. pp. 5589–5595 (May 2011). https://doi.org/10.1109/ICRA.2011.5980170
20. Lemaignan, S., Ros, R., Mösenlechner, L., Alami, R., Beetz, M.: Oro, a knowledge management platform for cognitive architectures in robotics. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 3548–3553 (2010). https://doi.org/10.1109/IROS.2010.5649547
21. Lim, G.H., Suh, I.H., Suh, H.: Ontology-Based Unified Robot Knowledge for Service Robots in Indoor Environments. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on **41**, 492–509 (Jun 2011). https://doi.org/10.1109/TSMCA.2010.2076404
22. Madin, J., Bowers, S., Schildhauer, M., Krivov, S., Pennington, D., Villa, F.: An ontology for describing and synthesizing ecological observation data. Ecological Informatics **2**(3), 279–296 (Oct 2007). https://doi.org/10.1016/j.ecoinf.2007.05.004
23. Masolo, C., Borgo, S.: Qualities in formal ontology (09 2010)
24. Mohammed Diab, Aliakbar Akbari, Muhayy Ud Din, Jan Rosell: PMK—A Knowledge Processing Framework for Autonomous Robotics Perception and Manipulation. https://www.mdpi.com/1424-8220/19/5/1166 (2019)
25. Noy, N.F., McGuinness, D.L.: Ontology Development 101: A Guide to Creating Your First Ontology (2001)
26. Prestes, E., Fiorini, S., Carbonera, J.: Core Ontology for Robotics and Automation (Sep 2014)
27. Redmon, J., Divvala, S.K., Girshick, R.B., Farhadi, A.: You only look once: Unified, real-time object detection. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016. pp. 779–788. IEEE Computer Society (2016). https://doi.org/10.1109/CVPR.2016.91, https://doi.org/10.1109/CVPR.2016.91
28. Rijgersberg, H., Wigham, M., Top, J.L.: How semantics can improve engineering processes: A case of units of measure and quantities. Advanced Engineering Informatics **25**(2), 276–287 (2011). https://doi.org/https://doi.org/10.1016/j.aei.2010.07.008
29. Rijgersberg, H., Van Assem, M., Top, J.: Ontology of units of measure and related concepts. Semantic Web **4**(1), 3–13 (2013). https://doi.org/10.3233/SW-2012-0069
30. Siciliano, B., Khatib, O. (eds.): Springer Handbook of Robotics. Springer (2008), http://dx.doi.org/10.1007/978-3-540-30301-5
31. Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A., Katz, Y.: Pellet: A practical OWL-DL reasoner. Journal of Web Semantics **5**(2), 51–53 (Jun 2007). https://doi.org/10.1016/j.websem.2007.03.004
32. Stenmark, M., Malec, J.: Knowledge-based instruction of manipulation tasks for industrial robotics. Robotics and Computer-Integrated Manufacturing **33**, 56–67 (Jun 2015). https://doi.org/10.1016/j.rcim.2014.07.004, citation Key: rosetta

33. Tenorth, M., Beetz, M.: KnowRob: A knowledge processing infrastructure for cognition-enabled robots. The International Journal of Robotics Research **32**(5), 566–590 (Apr 2013). https://doi.org/10.1177/0278364913481635
34. Thompson, A.: Guide for the Use of the International System of Units (SI)
35. Vrandečić, D., Krötzsch, M.: Wikidata: A free collaborative knowledgebase. Communications of the ACM **57**(10), 78–85 (Sep 2014). https://doi.org/10.1145/2629489
36. Waibel, M., Beetz, M., Civera, J., D'Andrea, R., Elfring, J., Gálvez-López, D., Häussermann, K., Janssen, R., Montiel, J., Perzylo, A., Schießle, B., Tenorth, M., Zweigle, O., De Molengraft, R.: RoboEarth. IEEE Robotics & Automation Magazine **18**(2), 69–82 (Jun 2011). https://doi.org/10.1109/MRA.2011.941632