


RESEARCH ARTICLE

Open Access



A development cycle for automated self-exploration of robot behaviors

Thomas M. Roehr^{1*} , Daniel Harnack¹, Hendrik Wöhrle^{1,2}, Felix Wiebe¹, Moritz Schilling¹, Oscar Lima¹, Malte Langosz¹, Shivesh Kumar¹, Sirko Straube¹ and Frank Kirchner^{1,3}

Abstract

In this paper we introduce Q-Rock, a development cycle for the automated self-exploration and qualification of robot behaviors. With Q-ROCK, we suggest a novel, integrative approach to automate robot development processes. Q-ROCK combines several machine learning and reasoning techniques to deal with the increasing complexity in the design of robotic systems. The Q-ROCK development cycle consists of three complementary processes: (1) automated exploration of capabilities that a given robotic hardware provides, (2) classification and semantic annotation of these capabilities to generate more complex behaviors, and (3) mapping between application requirements and available behaviors. These processes are based on a graph-based representation of a robot's structure, including hardware and software components. A central, scalable knowledge base enables collaboration of robot designers including mechanical, electrical and systems engineers, software developers and machine learning experts. In this paper we formalize Q-ROCK's integrative development cycle and highlight its benefits with a proof-of-concept implementation and a use case demonstration.

Keywords: Robotics, Self-exploration, Robot behaviors, Semantic annotation, Development cycle, Knowledge representation

Introduction

Modern robotics has evolved into a collaborative endeavor, where various scientific and engineering disciplines are combined to create impressive synergies. Due to this increasingly interdisciplinary nature and the progress in sensor and actuator technologies, as well as computing hardware and AI methods, the capabilities and possible behaviors of robotic systems improved significantly in recent years. Along with the greatly enhanced potential to strengthen established application fields for robotics and unlock new ones, these developments pose several challenges for developers and users interacting with robotic systems. On the one hand, for hardware and software engineers, these technological improvements led to an increasing size of the design space and, hence, development, integration and programming complexity.

Engineers do not only have to deal with technical peculiarities of a rich variety of different components when constructing a robot. They also have to develop advanced control strategies and integrate knowledge from a range of disciplines in order to unlock the full potential regarding a robot's capabilities.

On the other hand, the field of end users of applications for robotic systems widens, as more complex and versatile robots open up a wealth of novel applications for which robots were unsuited only years ago. In the last couple of years, the usage of robots has increased substantially and robots are used in more and more fields of application [1, 2]. These users will not be interested in the detailed construction of hardware or software, but will rather evaluate a robotic system by its possible behaviors and the tasks it can accomplish. However, without the domain knowledge of an experienced roboticist or AI researcher, designing a robot and the algorithms that provide the desired behaviors is next to impossible, and employing

*Correspondence: thomas.roehr@dfki.de

¹DFKI GmbH Robotics Innovation Center, Bremen, Germany
Full list of author information is available at the end of the article

engineers for construction of a custom robot is likely to be prohibitively expensive. Hence, especially small and medium-sized companies face difficulties to adopt robotic systems that suit their specific applications [3]. To overcome this problem, new robot engineering methods are required.

We claim that both collaborative teams of roboticists and end users would greatly benefit from a unifying automated framework for robot development that spans several abstraction levels to interact on. Our initial conceptual idea is outlined in [4], where we introduce Q-ROCK as a development framework leveraging integrative AI, which means that it applies and combines different AI technologies in one system to address the development change. With Q-ROCK we intend to simplify and automate the whole process from robot design to behavior acquisition and its final deployment for experienced roboticists from different disciplines and naive users alike. In the following sections we detail the implementation of the concept and explain essential elements.

The core hypothesis underlying this project is that the set of all possible behaviors of a robotic system is *inherently* defined by its constituting hard- and software, and, furthermore, that this set can be found by self-exploration of the robotic system. The major challenge is that the size of the behavior space is subject of the curse of dimensionality. In order to make the behavior space more manageable, we restrict the automated exploration to the kinematic capabilities of the robot. The Q-ROCK development cycle allows the usage of existing behavior components from other methods, e.g., for the interpretation of sensor input, thus permitting a complementary, semi-automated exploration approach. For the kinematic exploration, we make the assumption that the robot is explored in a minimal environment that allows to transfer explored capabilities to more complex environments. Additionally, we limit the capabilities to movements having a fixed length and we consider only kinematic states that the robot can reach through the movements itself. These simplifications still cannot negate the curse of dimensionality. At the lowest level, the robot, due to sensor resolutions and digital input signals, is a discrete system. Treated as such, a robot offers a number of possible states and transitions between states, which is not tractable for higher degrees of freedom in practical applications [5]. Therefore, in practice we are not able to simulate and store every possible movement individually. Instead, we construct a representative set of movements generated from a parameter space, which encodes all considered capabilities.

A distinct feature of our approach is that the self-exploration of the hardware is as goal-agnostic as possible, such that novel behaviors can be synthesized from already explored capabilities of the system without having to re-perform exploration with a novel task in mind. This

reusability is made possible by clustering capabilities and describing resulting behaviors in a semantically annotated latent feature space.

One important reuse of the capability clusters is the exploration of systems of systems. Here, the explored capabilities of the subsystems are used to efficiently explore the capabilities of the assembled new systems. Even though the exploration of simple subsystems can be already expensive, the general idea is to considerably increase sample efficiency for complex assembled systems. This paper deals with the exploration of base systems, which are required for the capability exploration of hierarchical systems. The latter, however, will be investigated in our future work.

Using a growing common knowledge base that links various description levels, from technical details of single components to behavior classification of self-explored robotic systems, we also provide a basis for behavior transfer between systems and reasoning about a possible robotic behavior given its composition. Hence, we propose a development cycle with multiple entry points that simplifies and speeds up the overall design process of robotic systems to benefit both developers and users.

Contributions

The main contribution of Q-ROCK is the integration of several different subdisciplines of AI into a common framework in order to explore and qualify robotic capabilities and behaviors. To this end, we integrate state-of-the-art methods and develop new approaches in four key areas:

- (i) Assembly of mechanical, electronic and software components with well defined interfaces and constraints
- (ii) Exploration and clustering of capabilities aided by machine learning
- (iii) Ontology-based semantic annotation of behaviors with user feedback
- (iv) Reasoning about a possible behavior given a robot's hard- and software composition

In this paper we introduce the formal concepts behind Q-ROCK and present a practical evaluation of our approach to tackle robot design problems in a combination of bottom-up and top-down solving. The evaluation is based on a *reach* behavior of a robotic manipulator, a mobile base, and a combination of both. Developing a simple behavior such as *reach* clearly does not justify a complex development cycle as presented in this paper. The implementation of this behavior, however, permits us to (a) outline a novel, integrative approach to develop robotic systems, which we refer to as the Q-Rock development cycle and (b) to provide a qualitative analysis of each

stage of this cycle. Hence, the analysis of the *reach* behavior serves to illustrate the essentials of the concept, and points to the potential of the Q-ROCK development cycle to achieve the automated exploration and classification of robot behavior.

The Q-ROCK development cycle supports and simplifies the interaction between different crafts in multiple ways. The traditional approach - with often strictly separated and sequenced steps of either hardware- or software-focused development - should be overcome. Firstly, by simplifying the robot design by means to reuse once designed components - where the link between hardware and software components is already established. Secondly, by an automated exploration of newly assembled systems, whose behaviors cannot always be inferred just from its components. We outline an approach that accelerates and simplifies robot design processes and permits robot developers and users to focus on optimization and fine-tuning.

Key elements of the Q-ROCK development cycle are: (a) a centralized data and knowledge base, (b) a procedure for automated system analysis by exploring a system assembly's possible behaviors (limited by some of our assumption), and (c) semantic annotation of these behaviors to enable their use in reasoning procedures. An important point to note is that Q-ROCK relies on a well-defined *robot hardware design process*, which is based on a data and knowledge base that provides the information to couple hardware and software components automatically by specifying requirements and compatibilities. We see this part as a crucial pre-requisite to implement the Q-ROCK concept, for which some foundations were developed in the precursor project D-ROCK [6]. Details of the robot hardware design process and its implementation are discussed in "[Modelling robot composition](#)" section.

As already stated, Q-ROCK combines different fields of research, where each area might have its own interpretation and definition for the same term. Where needed, in order to avoid confusion through conflicting connotations of a term, we decided to introduce a new one instead.

Paper outline

In the following "[State of the art](#)" section, we give a short overview of the current state and limitations of automatic robot behavior learning. "[Q-ROCK development cycle](#)" section introduces the concept of Q-ROCK, provides an overview of the methodology and formally defines the procedures and abstractions to implement Q-ROCK. "[Results: a use case scenario](#)" section describes exemplary use case scenarios to illustrate the implemented concepts presented in this paper. A discussion in "[Discussion](#)" section concludes the paper.

State of the art

Multiple disciplines, i.e., knowledge representation and reasoning, task planning as well as machine learning, can provide important methods to explore robotic capabilities or combine capabilities to generate more complex ones. However, to the best of our knowledge, there is little work in automated robot design that is approached in a holistic way as it is done in Q-ROCK. We will highlight relevant holistic approaches here, whereas related work in subdisciplines of the Q-ROCK development cycle are described in the corresponding paragraphs.

Ha et al. [7] suggest an automated method for the design of robotic devices using a database of robot components and a set of motion primitives. They use a high level motion specification in the form of end effector waypoints in task space. Their system then takes this motion specification as input and generates the simplest robot design that can execute this user-specified motion. However, Ha et al. do not consider the inverse problem which is the base concept of Q-ROCK: finding all motions a device can perform - under given assumptions.

A similar development, tackling the problem of learning motion behaviors via exploration is pursued in the project MEMMO (Memory of Motion) [8], where a graph in state space is generated during exploration, and where the links between nodes refer to control strategies adhering to the system dynamics. Both graph and control strategies are refined during exploration, and the resulting trajectories are then used during deployment to warm-start an optimal control framework. The key difference to our approach is that in the MEMMO framework, task objectives need to be known and encoded in a loss function for training, whereas our framework is mostly goal agnostic during exploration.

A system providing access to robotics development via a web based platform is included in the Amazon Web Services (AWS) [9]. The services include a RobotMaker which basically enables use of Robot Operating System (ROS) based tools via browser windows. This way the user doesn't have to install any tools locally. However, as far as it could be investigated, even though an account could be created freely, most of the services are commercial. Additionally, even though the ROS community provides many solutions for different applications, a tool that provides an easy access to a non-expert user, as aimed at by the Q-ROCK system, is lacking and is also not provided by AWS.

Another holistic approach for constructing and simulating robots is presented by the Neurorobotics Platform [10], under development within the Human Brain Project [11]. At the time of writing, this web-based framework includes an experiment designer, robot construction for simple toy robots (Tinkerbots [12]), a range of predefined robots and brain models, and various plotting and

visualization tools. The focus lies on fostering collaboration between neuroscientists and roboticists and providing simulated embodiment for biologically inspired brain models. In Q-ROCK we rather focus on exploration of possible capabilities given a robot's composition, and linking these capabilities and corresponding behaviors to its properties.

Q-Rock development cycle

To explore and annotate the inherent capabilities and possible behaviors of a robot and subsequently allow for reasoning about relations between composition and behaviors, Q-ROCK combines different kinds of AI techniques in a *development cycle* (see Fig. 1). This cycle can be driven by the high-level task specifications of a user, but is also flexible enough to support experienced domain experts. The cycle is divided into three major steps: (i) simulation-based exploration of the capabilities of a given piece of robot hardware, (ii) clustering and annotation of these capabilities to generate behaviors and behavior models, and (iii) model-based reasoning about the set of behaviors and related hard- and software that is required for a specific task. The Q-ROCK database - implemented as a combination of hand curated ontologies [13] and a graph database¹ - provides the central knowledge base to connect all steps. The database provides information about known hard- and software components and their relations, e.g., compatibility of component interfaces, and the structure of available robotic systems. As central storage of the results of each stage, the database enables the immediate use of data across all workflow steps, which leads to a fully integrated development workflow. The development cycle can be initiated from two entry points (E1 and E2 illustrated in Fig. 1). The first entry point E1 allows to enter a bottom-up development approach. Here, the goal is to identify the capabilities of a given robotic system or subsystem. E1 starts with hard- and software composition and ends up with all capabilities of that system, organized in semantically described cognitive cores.

The second entry point E2 represents a top-down approach. A user triggers the development cycle by providing a task definition, i.e., a given user scenario consisting of an environment and a specific problem that a robot, which is not known to the user, shall solve. The goal is to either find a robot in the database that is suitable to address the specified problem or to suggest a novel composition that will likely solve the task.

Complementary to the Q-ROCK development cycle overview in Fig. 1, we provide a standard Entity Relationship diagram in Fig. 2 to illustrate involved entities and their relationships. The following sections motivate and outline the different steps of the Q-ROCK development

cycle, and successively introduce these entities and their definitions to formalize our approach.

Modelling robot composition

For all steps of the development cycle it is essential to have a well-defined model of a robotic system. In Q-ROCK, we represent a robotic system, i.e., the specific types and compositions thereof - as well as relations between - robot hard- and software components, using a graph-based model.

Related work

The formal Architecture Analysis and Design Language (AADL) is designed to describe both processing and communication resources, as well as software components and their dependencies. A system designer is supposed to thoroughly model the system design, such that given an application designed by the application designer, it can be deployed on the system. Furthermore, it is possible to use special tools to perform design analysis prior to compilation and/or testing in order to find errors before deployment. A detailed overview of AADL can be found in [14].

TASTE is a framework developed by the European Space Agency to design, test and deploy safety-critical applications. It uses AADL as the modelling layer to design systems and applications. Based on these models the framework builds the glue code and enables the deployment of the software to a variety of different processing and communication infrastructures. Details can be found in [15].

In contrast to the aforementioned approaches, the domain-specific language NPC4 developed by Scioni et al. [16] uses hypergraphs to model all aspects of structure in system design, software design and other domains. Its four main concepts are node (N), port (P), connector (C) and container (C) combined with the two relations contains (C) and connect (C); refinements of these concepts form domain-specific sublanguages. A detailed description of the concept and the language NPC4 is presented in [16].

Our approach aims at exploiting the flexibility and formalization of NPC4 for a structural reasoning approach and combine it with the well-known and tested concepts of TASTE/AADL. However, unlike NPC4 our approach is based on standard graphs to make use of state-of-the-art database technology.

Approach & formalization

Components Components represent the hard- and software building blocks of robotic systems, which can be combined to generate more complex components. Hence, a hierarchy of components of different complexity is created. At the lowest level of this hierarchy are *atomic components*, which can not be divided into other components in our model.

¹<https://janusgraph.org>

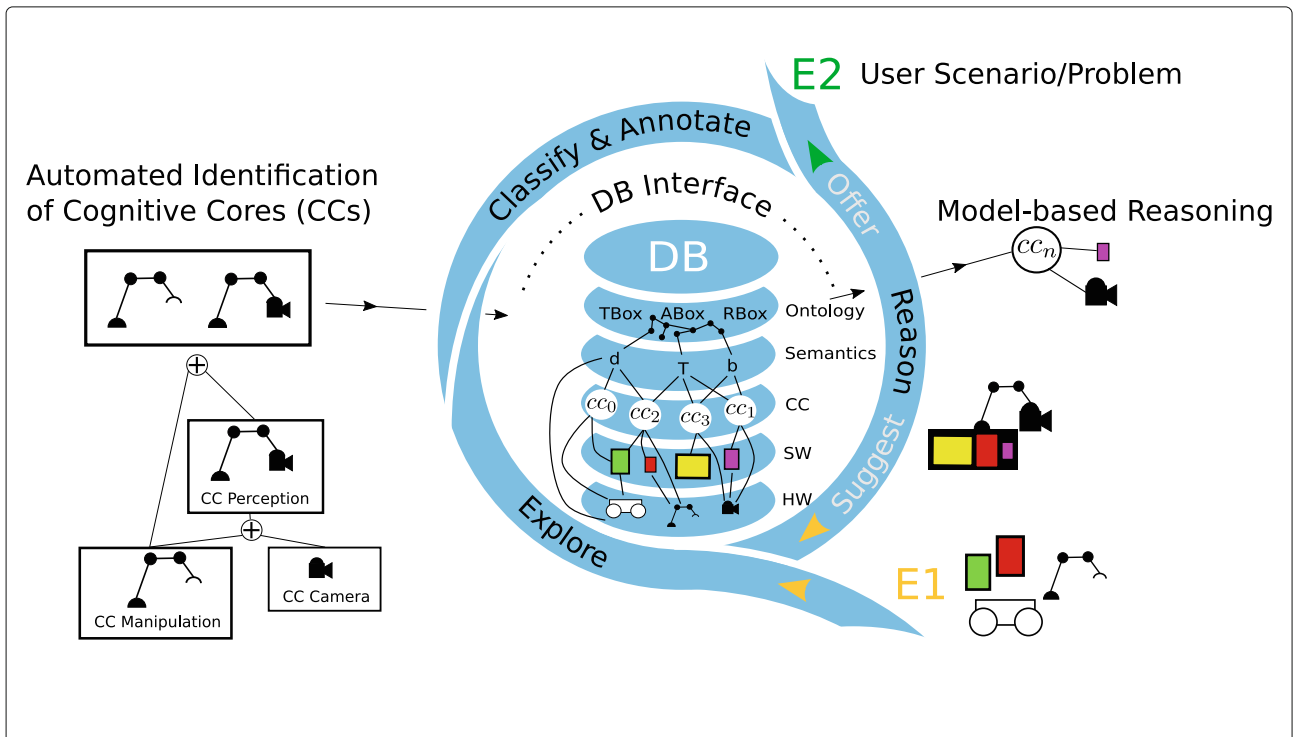


Fig. 1 The Q-Rock development cycle consists of three complementary steps: “Exploration”, “Classification and Annotation”, and “Reasoning”. A graph database serves as a central knowledge representation and data exchange hub. The process may be initiated from two entry points (E1 and E2), depending on the intention of a user. The entry point specifies whether Q-Rock follows a *bottom up* (E1) or *top down* (E2) development approach

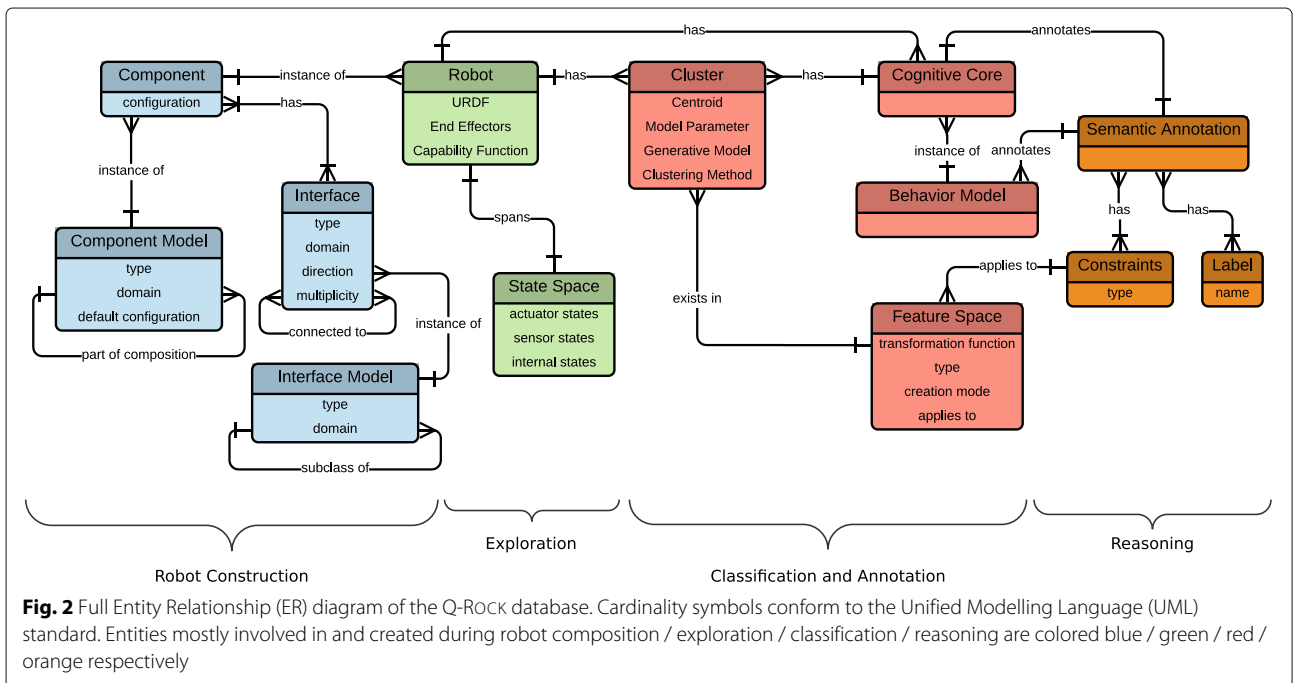


Fig. 2 Full Entity Relationship (ER) diagram of the Q-Rock database. Cardinality symbols conform to the Unified Modelling Language (UML) standard. Entities mostly involved in and created during robot composition / exploration / classification / reasoning are colored blue / green / red / orange respectively

Components are grouped into a predefined, but extendable set of domains $\mathcal{D} = \{\mathcal{S}, \mathcal{P}, \mathcal{M}, \mathcal{E}, \mathcal{A}\}$. The domains are described in Table 1. Each domain can only form new components by combining other components of the same domain (unless they are *atomic* components). The only exception is the *Assembly* domain which allows the composition of components of different domains. Thus, the *Assembly* domain is the one in which complete robotic systems - including their mechanical, electrical, processing and software structure - can be represented.

The main entities and their relationships are represented as labelled vertices and edges in a graph $G = (V, E, s, t, \Sigma, p_v, p_e)$. Here, V is the vertex set, E the edge set with s, t identifying source and target vertices, Σ is a vocabulary, and $K \subset \Sigma \setminus \{\emptyset\}$ is a set of predefined keys with 'label' $\in K$. Property functions for vertices and edges are defined as $p_v : V \rightarrow K \times \Sigma$ and $p_e : E \rightarrow K \times \Sigma$, where $p_v(\text{'label'}) \neq \emptyset$ and $p_e(\text{'label'}) \neq \emptyset$. The main entity sets are listed in Table 2, whereas relations are listed in Table 3. Note that all entities are represented as vertices in the graph, so that all entity sets listed in Table 2 are subsets of the vertex set V , and likewise all relations are subsets of the edge set E .

Relations have to be constrained to form a consistent system. The relations I_C, I_I, P, S are many-to-one relations; that means, that no element of their domain can be mapped to more than one element in their co-domain. This constraint ensures, for instance, that parts of one component model cannot be parts of another component model. The relations H, H_C are one-to-many relations, thus preventing ambiguity of interfaces between different entities. C_o is a many-to-many relationship, allowing any connection between interfaces, whereas the A relation is a one-to-one relation between an external and an internal

Table 1 Predefined domains for components in a robotic system

| Name | Symbol | Description |
|---------------|---------------|---|
| Software | \mathcal{S} | Software components which are, individually or in combination, used to interface and control a robotic system |
| Computational | \mathcal{P} | Physical entities which represent computational resources and the communication infrastructure |
| Mechanics | \mathcal{M} | Components required to establish a mechanical structure / kinematic design of a robotic system |
| Electronics | \mathcal{E} | Electronic devices including sensors, actuators, power supplies and supply lines |
| Assembly | \mathcal{A} | Composite components comprised of any set of components of the previously mentioned domains |

Table 2 Entities for modelling a robot's structure

| Name | Symbol | Description |
|-----------------|---------------|--|
| component model | M_C | All available component models in the respective domains |
| component | \mathcal{C} | Instances of component models |
| interface model | M_I | All interface types a component (model) could expose |
| interface | \mathcal{I} | Instances of interface models - possibly owned by some component (model) |

interface. Given the entities and their relations, the operators listed in Table 4 are defined on the graph. Algorithm 1 serves as example to illustrate the usage of these operators to construct the component model of a robotic leg. It is assumed, that a component model for a robotic joint $J \in M_C$ with two (external) mechanical interfaces $a, b \in \mathcal{I}$ exists. Furthermore, the existence of a component model for a robotic limb $L \in M_C$ with two (external) mechanical interfaces $x, y \in \mathcal{I}$ is assumed. Figure 3 visualizes the graph structure resulting from running Algorithm 1. The graph has three components g_i , gears of different ratio, in the mechanical domain $g_i \in \mathcal{M}$. One of it has been instantiated (I_C) and is part of (P) an actuator $A \in \mathcal{A} \cap M_C$. Chaining the respective relations $I_C^{-1} \circ P$ (see Table 3) resolves to:

$$\left\{ (g_i, a) \mid \exists x \in \mathcal{C} : (g_i, x) \in I_C^{-1} \wedge (x, a) \in P \right\} \text{ for some } i.$$

The component (instance) actuator a with stator and gear as its composing parts is combined with controller electronics and controller software to define the joint model. This model defines the structure of joint instances in the higher-level leg component.

Algorithm 1 Example application of graph operators to constructing a component model

```

j1 ← instantiate_C(J, hip1)
j2 ← instantiate_C(J, hip2)
j3 ← instantiate_C(J, hip3)
j4 ← instantiate_C(J, knee)
connect(j1.b, j2.a)
connect(j2.b, j3.a)
l1 ← instantiate_C(L, upperLimb)
l2 ← instantiate_C(L, lowerLimb)
connect(j3.b, l1.x)
connect(l1.y, j4.a)
connect(j4.b, l2.x)
m = create_C(A, Leg)
compose(j1, m)
compose(j2, m)
    
```

Table 3 Available relation types to describe relationships between entities

| Relation name | Definition | Cardinality | Description |
|-------------------------------|---|-------------|--|
| instance of (component model) | $l_C : \mathcal{C} \rightarrow M_C$ | N:1 | specify the model $x \in M_C$ a component $y \in \mathcal{C}$ has to comply with |
| instance of (interface model) | $l_I : \mathcal{I} \rightarrow M_I$ | N:1 | specify the model $x \in M_I$ an interface $y \in \mathcal{I}$ has to comply with |
| subclass of | $S : M_C \rightarrow M_C$ | N:1 | establish a hierarchy among the component models |
| (model) has interface | $H : M_C \rightarrow \mathcal{I}$ | 1:N | define the interfaces a component model exposes to the exterior |
| (component) has interface | $H_C : \mathcal{C} \rightarrow \mathcal{I}$ | 1:N | define the interfaces a component exposes to the exterior, given the model of it as a template |
| part of composition | $P : \mathcal{C} \rightarrow M_C$ | N:1 | partially define the inner composition of components of other components |
| connected to | $C_o : \mathcal{I} \rightarrow \mathcal{I}$ | M:N | connect components to form networks, thus refining the inner composition |
| alias of | $A : \mathcal{I} \rightarrow \mathcal{I}$ | 1:1 | link interfaces in the interior to the exterior of a component model |

Implementation

The robot hardware design process presented in this section has been implemented using a state-of-the-art graph database. The database directly supports the creation of a graph G with labelled vertices and edges and - most notably - the direct specification of cardinality constraints depicted in Table 3. The database is accessed at low-level via the *Gremlin query language* [17] and a higher-level through a Python layer that we implemented. The higher-level Python layer translates (sub-)graphs of the database into a network of interconnected Python objects. It is also this layer in which the graph operators of Table 4 are implemented as the basis for the whole Q-ROCK toolchain.

Exploration

The exploration step in the Q-ROCK development cycle aims at finding all capabilities a given piece of robotic hard- and software can provide in a meaningfully defined

exploration environment. The chosen exploration environment has to allow a transfer to more complex application environments. Capabilities in this context mean the possible trajectories a robot can produce within a robot state space and a world state space. The exploration is based on simulating a robot that has been modelled as formally described in “[Modelling robot composition](#)” section - a practical example is given in “[Results: a use case scenario](#)” section.


Related work

In distinction to other approaches in robotics, we try to avoid directing the exploration towards any kind of goal, and instead aim at generating a maximal variety of capabilities to find a representative set that might include novel, unanticipated ones.

Capability exploration methods usually aim to create a library of diverse capabilities [18–21], so that the coverage of the behavior space is maximized and the capabilities

Table 4 List of available operators, which allow to modify the graph structure

| Name | Definition | Description |
|--------------------------------|--|--|
| <i>create_C</i> | $\mathcal{D} \times \Sigma \rightarrow M_C$ | add a component model to the graph |
| <i>create_I</i> | $\mathcal{D} \times \Sigma \rightarrow M_I$ | add an interface model to the graph |
| <i>instantiate_C</i> | $M_C \times \Sigma \rightarrow l_C$ | create a component c from a model m s.t. $(c, m) \in \mathcal{I}_C$ |
| <i>instantiate_I</i> | $M_I \times \Sigma \rightarrow l_I$ | create an interface i from an interface model m s.t. $(i, m) \in l_I$ |
| <i>isA</i> | $M_C \times M_C \rightarrow S$ | make component model $x \in M_C$ a subclass of component model $y \in M_C$, such that $(x, y) \in S$ |
| <i>has_M</i> | $M_C \times \mathcal{I} \rightarrow H$ | associate an interface instance with a component model |
| <i>has</i> | $\mathcal{C} \times \mathcal{I} \rightarrow H_C$ | associate an interface instance with a component |
| <i>compose</i> | $\mathcal{C} \times M_C \rightarrow P$ | make a component part of a component model |
| <i>connect</i> | $\mathcal{I} \times \mathcal{I} \rightarrow C_o$ | define a connection between two interface instances |
| <i>export</i> | $\mathcal{I} \times \mathcal{I} \rightarrow A$ | define one interface to be an alias for another interface, e.g., to map a component model’s interface to a composing component’s interface |



can be utilized in different tasks and environments. For the capabilities to be transferable between tasks, these approaches avoid task specific reward functions. Instead they use intrinsic motivations such as novelty, prediction error and empowerment. An extensive overview of intrinsic motivations in reinforcement learning can be found in [22].

Because of their inherent incentive to explore and find niches, evolutionary algorithms are natural candidates for behavior exploration. Lehman and Stanley [23] propose to use *novelty* as the sole objective of likewise-called novelty search. It was found to perform significantly better than goal-oriented objectives in deceptive maze worlds. Novelty search has already been applied to robotics to find multiple *diverse* high quality solutions for a single task [24, 25]. Cully [26] suggests combinations of different methods, e.g., quality-diversity optimization methods and unsupervised methods, which allow to explore various capabilities of a system without any prior knowledge about their morphology and environment.

Formalization

The abstracted approach of the exploration is depicted in Fig. 4. It serves as high level description of the process where the formalization will be given in the coming paragraphs. Exploration discovers a set of capabilities by applying a search strategy, where the challenge lies in handling a significantly large state space. We tackle the large state space using a parameter-only based encoding

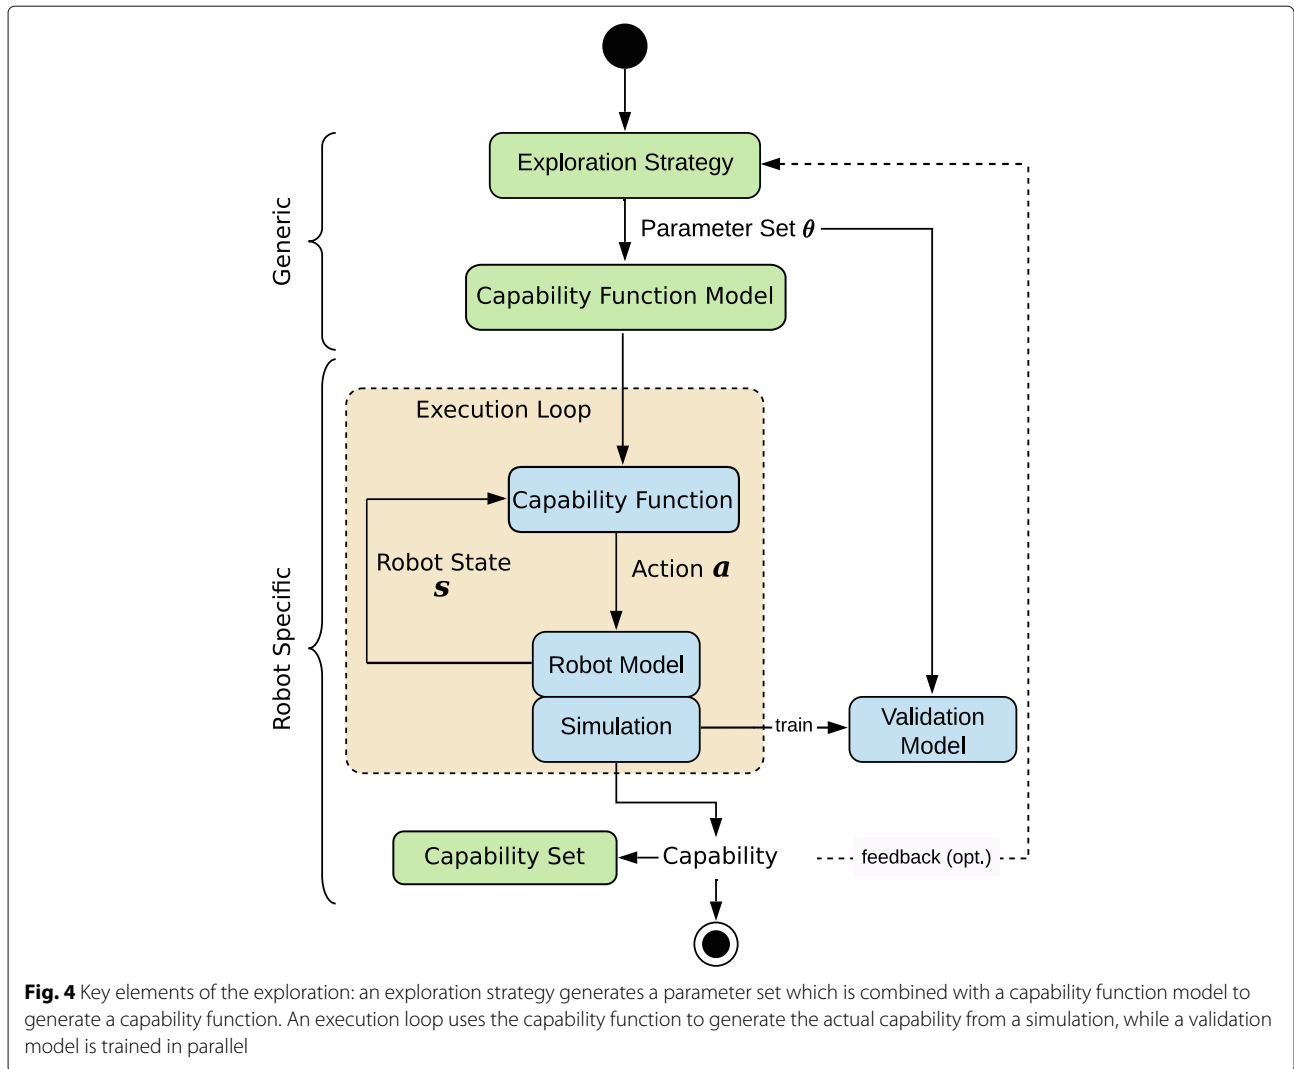
for capability functions: the encoding is compact and yet arbitrarily precise. Creating a capability function from a dedicated capability function model and applying it on the actual robot in an execution loop results in a capability of the system, where a capability is the executed trajectory in the world state space. This structure allows to validate the feasibility and cluster robot-specific execution characteristics from capabilities on the basis of the input parameter space.

States

Definition (Joint State). *The joint state $q \in Q$ is a vector of all joint positions of the robot. For a robot with n joints: $Q \subset Q^1 \times Q^2 \times \dots \times Q^n$. Q is a subset of the Cartesian product because not all combinations of joint positions may be allowed due to the robot structure.*

Definition (Actuator State). *The actuator state s^a of the robot is a tuple of the configuration and the joint velocity, so that $s^a = (q, \dot{q}) \in S^a$, where the actuator state space $S^a = Q \times V_Q$ combines joint state space Q with joint velocity space V_Q . The robot actuator state completely describes the positions and velocities of all parts of the robot at a given time.*

The complete (observable) state of the robot contains not only the actuator state s^a but also the states $s^s = (s^{s,1}, \dots, s^{s,m}) \in S^s$ of all m sensors and possibly internal states s^i .



Definition (Robot State). The (full) robot state \mathbf{s}^{rob} is a combination of actuator, sensor and internal states. An internal robotic state $\mathbf{s}^i = (s^{i,1}, \dots, s^{i,k})$ for k internal properties and $\mathbf{s}^i \in \mathcal{S}^i$ may encompass for example internal time, battery status or a map of the robot's surroundings. $\mathcal{S}^a, \mathcal{S}^s$ and \mathcal{S}^i are the sets of all possible actuator, sensor and internal states, respectively. The full robot state reads:

$$\mathbf{s}^{rob} = (\mathbf{s}^a, \mathbf{s}^s, \mathbf{s}^i) \in \mathcal{S}^{rob} = \mathcal{S}^a \times \mathcal{S}^s \times \mathcal{S}^i.$$

The robot state \mathbf{s}^{rob} does not contain the complete information about the actual physical or internal state of the robot. It does only contain information that is accessible by the robot itself, i.e., that can be captured. Information about sensorless unactuated joints for example is not part of the robot state.

Actions To trigger changes of the robot state and thereby generate trajectories of robot states, it is necessary to

define motor actions. A motor, which is part of a joint, outputs a motor torque $\tau \in \mathcal{T}$. The torques for all joints can be written as a tuple $\boldsymbol{\tau} \in \mathcal{T} = \mathcal{T}^1 \times \mathcal{T}^2 \dots \mathcal{T}^n$. An idle joint always outputs $\tau = 0$.

Definition (Action). A kinematic action \mathbf{a}^{kin} is a tuple of a torque $\boldsymbol{\tau} \in \mathcal{T}$ and a time interval Δt , such that $\mathbf{a}^{kin} = (\boldsymbol{\tau}, \Delta t) \in \mathcal{A}^{kin}$, where \mathcal{A}^{kin} denotes the kinematic action space. Applying a kinematic action to the robot maps the current robot state to a new robot state:

$$\mathbf{a}^{kin} : \mathcal{S}^{rob} \rightarrow \mathcal{S}^{rob}$$

Besides kinematic actions, there are also perceptive actions $\mathbf{a}^{per} \in \mathcal{A}^{per}$ which evaluate sensor data and store abstractions in the internal robot state:

$$\mathbf{a}^{per} : \mathcal{S}^s \rightarrow \mathcal{S}^i$$

Finally, there are internal actions $\mathbf{a}^{int} \in \mathbf{A}^{int}$ processing internal information:

$$\mathbf{a}^{int} : \mathbf{S}^i \rightarrow \mathbf{S}^i$$

The full action space is the Cartesian product of the individual action spaces:

$$\mathbf{A} = \mathbf{A}^{kin} \times \mathbf{A}^{per} \times \mathbf{A}^{int}$$

Environments & world state Note that environments with different properties, including, but not limited to, gravitational force, pressure, and temperature will have an influence on the outcome of a kinematic action. Hence, environmental parameters as well as poses and properties of objects in the robot's workspace have to be considered when evaluating the feasibility of kinematic actions.

Furthermore, the environment is an important component to identify certain properties of capabilities. For example a throwing capability relies on the temporal evolution of states of the object to be thrown, which is represented by environment states that can be external to the robot (if it does not have the appropriate sensing capabilities). This arises also for capabilities that can, at first glance, be considered mostly environment independent: The effect of actions on the trajectory of an end effector when pointing is still determined by gravity and the viscosity of the medium in which the movement is performed. Even more, there is no generic way to determine the poses of all the robot's limbs just from sensing the actuator states \mathbf{s}^a : if a system is underactuated or does not have sensors on some actuators, an analytical solution for the feed forward kinematics may not exist. To compensate for this, we introduce the world state space which may also contain information unavailable to the robot itself:

Definition (World State Space). *The world state $\mathbf{s}^{world} \in \mathbf{S}^{world} = \mathbf{S}^{rob} \times \mathbf{S}^{obs}$, where the observational state space \mathbf{S}^{obs} contains states read from the environment, e.g., the position and orientation of objects or robot limbs and end effectors. These states are obtained during simulation or by monitoring a real world execution and will be accessible to the robot if it has the appropriate sensing capabilities.*

Capability A particular capability will require the sequential execution of a sequence of actions. Such an action sequence can be represented by a capability function. The capability function selects an action for the robot based on the current state and time, and thus defines how the robot is supposed to (re-)act in a given situation.

Definition (Capability Function). *A capability function is a function cap that maps the robot state at a given time to an action:*

$$cap : \mathbf{S}^{rob} \times \{0, \dots, T\} \rightarrow \mathbf{A}$$

and $cap \in \mathbf{CF}$, where \mathbf{CF} denotes the capability function space. An important detail to note is that the capability function operates on the robot state space and not the world state space. A capability function is a robot inherent function that considers only information that is available to the robot itself.

A capability function can be created in various ways, e.g., it could be a policy obtained from reinforcement learning, a behavior from an evolutionary algorithm, or a control law from optimal control theory.

In general the generation of capability functions can be formulated with a capability function model:

Definition (Capability Function Model). *A capability function model is a mapping cfm from a parameter space Θ to the capability function space*

$$cfm : \Theta \rightarrow \mathbf{CF}$$

The capability function model introduces a parameter space Θ , which allows the parametric generation of capability functions and is the basis for the exploration. In order to not constrain the exploration, the capability function model should be able to represent all kinds of capability functions of a system. In principle, however, it is also possible to operate with multiple capability function models at once.

By repeatedly calling a capability function and applying the resulting actions to the robot, a capability is executed.

Definition (Capability). *A capability $\mathbf{l}^T \in \mathbf{L}$, where T is the finite time horizon and \mathbf{L} the space of all trajectories, is defined by a sequence of world states and time coordinates:*

$$\mathbf{l}^T = \left[(\mathbf{s}^0, t^0), (\mathbf{s}^1, t^1), \dots, (\mathbf{s}^T, t^T) \right]$$

where the transition between successive states \mathbf{s}^t and \mathbf{s}^{t+1} is effected by an action \mathbf{a}^t of the robot.

Capabilities are central entities in the Q-ROCK philosophy, since we argue that a complete set of all possible capabilities is the most fundamental representation of what a system is able to do. In general the result of executing a capability strongly depends not only on the robot itself but also on the environment. With the intention to isolate the capabilities of the robot itself, we assume the environment to be minimalistic, deterministic, and static.

To refine the notion of completeness, we make two crucial assumptions at this point:

Assumption (Discretized Time). *We assume a discretized time model by arguing that (most) robots are controlled by digital hardware or controllers which have a specific clock or controller frequency. The smallest time step considered here is the denoted by δt .*

Assumption (Discretized State Space). *We assume that \mathcal{S} is sensed by digital sensors and we consider only state changes if we can distinguish them. As consequence we have a discretized state space \mathcal{S} .*

While this discretization reduces the cardinality of \mathcal{L} , it is still countably infinite if t is not bounded, so we have to choose a maximal capability length T . Now, in principle, a complete set of all possible capabilities up to a maximal length can be generated. Not surprisingly, this set would still have an intractable size considering typical resolutions of modern hardware and degrees of freedom [5]. A possible approach is to use a generic capability representation instead of manually specifying a sequence of actions. The parameters of such a representation define the resulting capability. Possible tools are motion primitives such as polynomials used by Kumar et al. [27], DMPs used by Schaal [28] or Gaussian kernel functions used by Langosz [29].

As a full set of capabilities is not tractable, the next best thing is a representative set of capabilities with a uniform distribution in a given feature space. An evolutionary algorithm such as novelty search [23] offers a suitable approach. With novelty search it is possible to search for novel capabilities with respect to a previously specified characteristic. A possible problem for goal-agnostic exploration however is precisely this characteristic that is used to define novelty, since it can bias the distribution of capabilities in an undesirable way for subsequent clustering, and partially preempts the distinct feature generation step. An alternative strategy is simple random sampling on the parameter space Θ , which also comes with its own caveat: The final distribution of the capabilities in later defined feature spaces will depend on the specific parametrization. Despite this, we see both approaches as viable alternatives, since they pose relatively low restrictions on the feature generation compared to more goal directed exploration. A representative set of capabilities, obtained with an exploration strategy like this, may serve as a starting point for exploring the space in a finer resolution, for capturing the system dynamics in a model, or for searching for a specific capability.

Because the capability function model itself is robot-agnostic, it is a priori not clear, which parameters θ correspond to feasible capabilities of the robot. For this reason, a validation model is trained that predicts which parameters θ lead to capabilities that are actuatable on the robot in the current environment.

After the exploration phase is finished, the obtained library of capabilities, the capability function model with the simulator to generate new capabilities, and the validation model are saved to the database and can be used by the following step.

Classification and annotation

The goal of this workflow step is the creation of cognitive cores. Cognitive cores are hubs that connect a specific behavior model with the robot's hard- and software, semantic annotations of that behavior model, and robot-specific capabilities that execute the behavior. Cognitive cores allow the execution of the corresponding behavior by using constraints and target values in semantically annotated feature spaces, and rely on clustering of capabilities in these spaces. Cognitive cores are central entities in Q-ROCK since they constitute our solution to the symbol grounding problem, i.e., link semantic descriptions to sub-symbolic representations, and serve as a basis for reasoning about the relation of hard- and software components, robotic structure, and resulting behavior.

Related work

One important point of our approach is the clustering of capabilities into feature spaces and the control of the robot within these feature spaces. Several studies have shown performance and robustness benefits from controlling a simulated agent in a latent, compressed feature space. Ha et al. [30] used a variational autoencoder on visual input to control a car in a 2D game world. In the context of hierarchical reinforcement learning, Haarnoja et al. [31] showed that control policies on latent features outperform state-of-the-art reinforcement learning algorithms on a range of benchmarks, and Florensa et al. [32] found high reusability of simple policies spanning a latent space for complex tasks. In a similar vein, Lynch et al. [33] investigate using a database of play motions, i.e., teleoperation data of humans interacting in a simulated environment from intrinsic motivation, combined with projections into a latent planning space to generate versatile control strategies. While the latter study has parallels to our segmentation into an exploration and a clustering phase, no previous approach aims at a semantically accessible feature representation, as we propose in Q-ROCK.

A method for generation of a disentangled latent feature space from observations was developed by Higgins et al. [34]. This variational autoencoder builds on the classical autoencoder architecture [35, 36] that compresses data into a latent space. The authors note that disentangling seems to produce features that are also meaningful in a semantic sense, such that changes in feature space lead to interpretable changes in state space.

Close to our work, Chen et al. [37] use a combination of variational autoencoders and Dynamic Movement Primitives (DMPs) to learn and generate robotic motion. They showcase that semantically meaningful representations of motion can be obtained, and that switching between different motions can be performed smoothly. To a similar end, Wang et al. [38] combine a variational autoencoder with generative adversarial imitation learning and show that a semantic embedding space can be learned for reaching and locomotion behavior. Whereas both these approaches highlight the construction of a semantically interpretable latent space for motion behaviors, they rely on human training data with a known semantical context. In our framework we aim at extracting the semantics after a motion library has been generated during the exploration step, which has no notion of semantics per se.

A combination of unsupervised clustering and variational autoencoders is described by Dilokthanakul et al. [39]. However, direct semantic annotation of these features and a formalized combination into behavior models has not been considered to date.

Formalization

To arrive at a formal definition of cognitive cores, we first need to clarify what constitutes a behavior. Since the term “behavior” has overloaded definitions in various disciplines, we specifically mean behavior in a broad, radical behaviorist sense, while emphasizing the phenomenological aspects: Everything an agent does is a behavior, and all behaviors must be in principle completely observable [40]. The complete observation is provided by the capabilities as defined in “Exploration”. We further define a behavior model as an abstraction of similar capabilities that have the same semantic meaning: A behavior of “walking” is not bound to the exact execution of a sequence of robot and world states, but rather a large number of capabilities that can differ in certain aspects. We thus propose that different behavior models can be identified by finding constraints to capabilities in appropriate feature spaces, leading to the following definition:

Definition (Behavior Model). *A robot-agnostic, semantically labelled abstraction of a set of capabilities L that adhere to constraints in feature spaces.*

Feature spaces arise from transformations of the capabilities via a feature function to capture specific aspects, and allow to define distances between capabilities within these aspects:

Definition (Feature Function). *A feature function ff_k maps from capabilities $l \in L$ to a set of values in \mathcal{R}_k^n , so that $ff_k : L \rightarrow \mathcal{R}_k^n$. This function is supplemented with a semantic description.*

Definition (Feature Space). *A metric space F_k with elements $f_k = ff_k(l)$. It is uniquely defined by the combination of ff_k and its metric m_k .*

An important aspect of the feature functions is their semantic descriptions, which constitute the language in which the behavior models are defined.

The feature functions ff_k can be obtained in two different ways. Either they are defined manually and directly annotated by a semantic description, e.g.,

$$ff_k(l^T) = \frac{1}{T+1} \sum_{t=0}^T \dot{q}^t$$

with the description ‘average actuator velocity’. Alternatively, they can be found automatically in a purely data-driven way, e.g., by using variational autoencoders [34] adapted to trajectory data. The Q-ROCK framework allows both approaches that can also be used in parallel to provide maximal flexibility, whereas the latter is not implemented to date. We thus enable the use of expert knowledge to define the most relevant feature spaces for a given problem. However, a non-expert user could also solely resort to the automatic approach. In addition, interesting feature functions that reflect the specifications of the robot model might be discovered automatically that are not obvious – even to an experienced observer, or hard to formulate.

An example of a behavior model defined by constraints in feature spaces is:

```
label : reach
constraints :  $F_1$  : min : 0.95   max : 1.0
               $F_2$  : variable
               $F_3$  : variable
```

Currently, only min/max and variable constraints are implemented. A variable constraint means that a target value f_k^{tar} has to be provided when the associated behavior should be executed by a robot. Since behavior models are robot-agnostic, they can be grounded for different robotic systems. The robot-agnostic nature of the behavior model depends on feature functions’ semantic descriptions: feature functions having the same effect on a semantic level may have varying definitions for different robots, especially if they are represented by encoder networks or other function approximators. Thus it must be possible to identify feature functions across robots by their semantic description.

To achieve a robot-specific grounding of the behavior model, the feature spaces F_0, \dots, F_k are populated by mapping robot-specific capabilities provided by the exploration step via the associated feature functions ff_1, \dots, ff_k .

In principle, if all capabilities of a robot are contained in the representative set provided by the exploration step,

a simple lookup of capabilities that adhere to the behavior model constraints is sufficient to execute the desired behavior. However, as noted before, this usually implies a capability set of intractable size.

We tackle this problem in two ways: Firstly, the capabilities are clustered in $F_k \forall k$ and the centroids of the clusters are used to check constraints for all members of the cluster. Whereas the result of this check is not exact for all capabilities, computational performance is greatly increased. Secondly, to avoid a lookup search when executing a behavior and to not be restricted to capabilities seen during exploration, we abstract generative models on the parameter sets θ from the capability clusters. Thus, clusters are represented by probabilistic generative models that, when sampled from, provide parameters θ which, via recurrent execution of the capability function model, lead to capabilities that likely lie in the intended clusters. Clusters are thus defined as:

Definition (Cluster). A cluster with label c_k^j is defined within a feature space F_k , which is associated with several clusters $j \in [1, n_k]$, where n_k denotes the number of clusters found in F_k . Each cluster has a generative model $G_k^j(\theta) \approx p(\theta, c_k^j) = p(\theta | c_k^j) p(c_k^j)$, that represents a probability distribution over parameter space θ , and a

centroid $f_k^j = \text{ff}_k \left(l \left(\arg \max_{\theta} G_k^j(\theta) \right) \right)$, where we use $l(\theta)$ as a shorthand for the combination of capability function model and recursive application of the execution loop (see Fig. 4).

Using generative models has the advantage that models from different clusters can be combined and jointly optimized to find a parameter set θ that generates a capability lying in several intended clusters. The clustering procedure is visualized in Fig. 5.

After clustering, robot-specific cognitive cores can be instantiated. Cognitive cores are defined as:

Definition (Cognitive Core). A cognitive core is an executable grounding of a behavior model for a specific robotic system, where constraints of the behavior model are checked against cluster centroids f_k^j . Clusters that satisfy these constraints are linked to the cognitive core. A cognitive core can only be generated when all behavior model constraints can be met.

These cognitive cores are described by a semantic annotation:

Definition (Semantic Annotation). A tuple $SA = (L, \mathcal{X})$, where L is a set of labels, $|L| \geq 1$ and \mathcal{X} is the set of constraints.

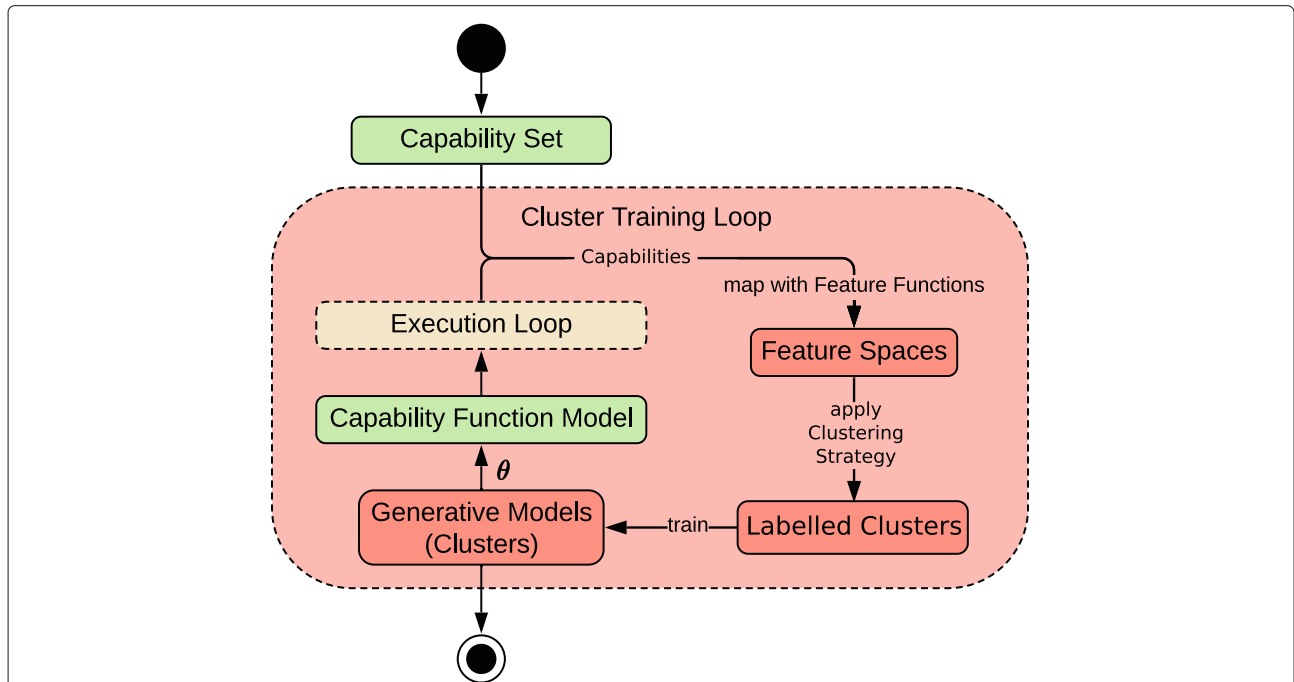


Fig. 5 Clustering overview. A representative capability set, along with the corresponding parameters θ and capability function model is provided by the exploration. Transformation functions ff_k are applied to map to feature spaces F_k . In these feature spaces, clustering is performed. The labelled clusters are used to train probabilistic generative models on the parameter space Θ , s.t. clusters can be stored in an efficient and expressive way. When sampling from the generative cluster models, parameters θ are generated that lead to capabilities in the intended cluster. The mapping from parameters to a capability is mediated by the capability function model and the execution loop (see Fig. 4). During training, sampling of parameters and generation of new capabilities is used to verify model performance

By default, the cognitive core inherits the labels and constraints from its behavior model, but the semantic annotation can be augmented by robot-specific information. This semantic annotation is the main interface between the generation of cognitive cores and the reasoning processes described in section “Reasoning”. The relation between feature spaces, clusters, behavior models, cognitive cores and semantic annotations is illustrated in Fig. 6.

In this framework, the execution of a behavior on a specific robot, i.e., the execution of a cognitive core, comes down to finding a parameter set θ_{max} that jointly maximizes all generative cluster models adhering to the constraints of the behavior model. If a behavior model includes variable constraints, each target value f_k^{tar} in the corresponding feature space F_k needs to be assigned. The cognitive core then finds the cluster models with closest centroids to the variable inputs. The cognitive core

effectively uses a constraint checking function $cc(c_k^j)$ to determine the relevant clusters, where

$$cc(c_k^j) = \begin{cases} 1 & \text{if type is "min/max" and} \\ & \min < \bar{f}_k^j < \max \\ 1 & \text{if type is "variable" and} \\ & \bar{f}_k^j = \arg \min_{\bar{f}_k^i} m_k(\bar{f}_k^i, f_k^{tar}) \\ 0 & \text{else.} \end{cases}$$

with function m_k as the metric of the feature space F_k . Note that this implies that several cluster models in the same feature space can fulfill a min/max constraint. The product of all currently relevant cluster models, i.e., the models G_k^j for which $cc(c_k^j) = 1$, results in a new probability distribution. The maximum of this distribution corresponds to a parameter set θ_{max} that has the highest likelihood of generating a capability that lies within

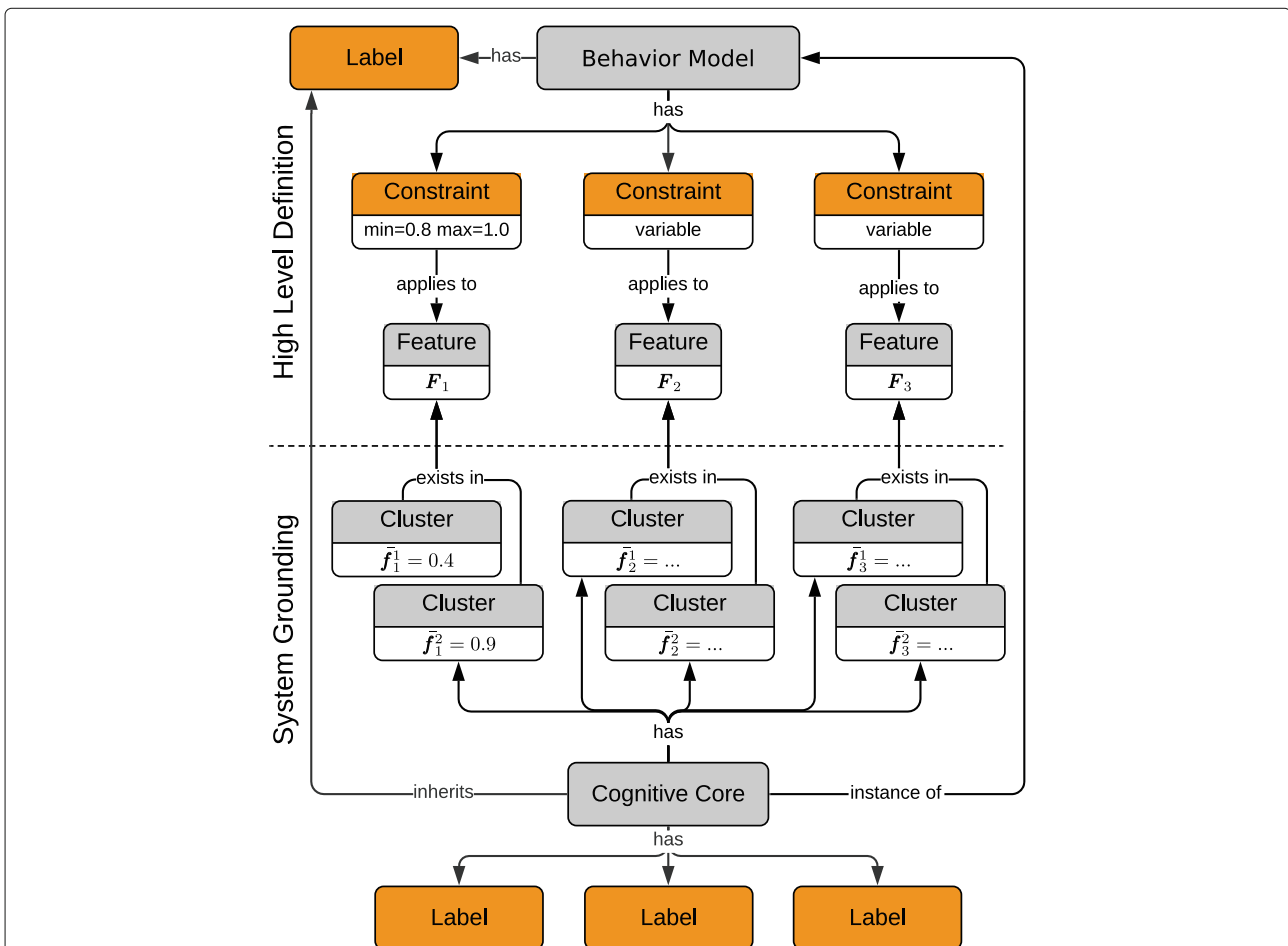


Fig. 6 Relation between features, clusters, behavior models, cognitive cores and semantic annotations. The behavior model is defined by constraints in feature spaces. This behavior can be grounded as a cognitive core for a specific system when clusters for this system exist that fulfill these constraints. Constraint fulfilling clusters are linked to the cognitive core. The cognitive core inherits the generic label of the behavior model, but can have more that describe specifics for this robot. The constituents of the semantic annotation are colored orange

all relevant clusters when used as input to the capability function model and the execution loop (see Fig. 4). The maximization step is then formally written as:

$$\theta_{max} = \arg \max_{\theta} \prod_{\mathbf{M}} G_k^j(\theta),$$

where $(k, j) \in \mathbf{M}$ if $cc(c_k^j) = 1$. Since this approach is based on probabilistic modeling, it is possible that the capability associated with θ_{max} violates a constraint. However, assuming a smooth mapping $\Theta \rightarrow \mathbf{F}_k$ via the feature functions ff_k , the violation is likely mild. If not violating a particular constraint is important, e.g., to avoid collisions, different weights can be assigned to different constraints, which control the relative influence of the corresponding cluster models. Note that it is also possible that cluster models are combined that have close to or completely disjunct distributions. Thus, in practice a probability boundary has to be set under which the maximization result θ_{max} is rejected and it is assumed that no capability exists that fulfills all constraints.

One important challenge of the approach is how behaviors are cast into the constraint-based, phenomenological behavior model we use. Since we aim at semantics which is intuitively understandable, we rely on human interaction. Thus, the first option Q-ROCK provides is hand-crafting behavior models. Although it requires some domain knowledge, this approach scales well in the sense that once defined, the behavior model can be grounded for many different robotic systems. In addition, we also envision semi-automated approaches: (1) Behavior modelling from observation of human examples, and (2) Modelling human evaluation functions with respect to a specific behavior. Approach (1) is based on research on end effector velocity characteristics for deliberate human movement [41, 42]. These movement characteristics can be formulated as feature space constraints and thus used to define behavior models. For approach (2), it was shown that implicit bio-signals of the human brain and explicit evaluation of a human observing simulated robot behavior can be used to effectively train a model of the underlying evaluation function [43], and to guide a robotic learning agent [44]. Also here, feature space constraints can be derived from the trained evaluation function approximator and used to define the behavior model.

At this point, we want to stress again that human interaction is absolutely necessary in the Q-ROCK philosophy to define meaningful behavior. Throughout this workflow step, human labelling is required for feature spaces, cognitive cores and behavior models. The robot itself, after exploration, has no notion of causality, i.e., reaction to the environment, or purpose in what it is doing. Thus it is not behaving in the actual sense. Only through human semantic descriptions, i.e., what it would look like if the

robot would behave in a certain way, are the capabilities of the robot in the environment ascribed to a meaningful behavior. Once the Q-ROCK database grows, we will explore automatically generated labelling of feature spaces based on similarity to already labelled ones, which could speed up the labelling process by providing reasonable first guesses.

To summarize, cognitive cores derived from behavior models are central entities in the Q-ROCK workflow, since they cast explored capabilities in a semantically meaningful form and provide a way to generate new capabilities that adhere to characteristics found by clustering. In addition, their semantic annotation provides the basis for reasoning about the connection of possible robot behavior to the underlying hard- and software, which will be elaborated in the following section.

Reasoning

Structural reasoning serves two purposes in Q-ROCK: (1) to suggest suitable hardware to solve a user-defined problem, and (2) to map an assembly of hardware and software to its function. The former does not involve any type of active usage of the hardware and software assembly, but exploits knowledge about the physical structure, interface types and known limitations / constraints when combining components, as well as their relation to labelled cognitive cores.

Essentially, structural reasoning establishes a bi-directional mapping between assemblies of hardware and software components and its function. Note, that we explicitly do not use the term robot here, since the result of the mapping from capabilities might not be a single robot, but a list of hardware and software components.

Related work

Knowledge Representation and Reasoning (KR&R) is considered a mature field of research, but there is still a gap between available encyclopedic knowledge and robotics. KNOWROB [45], as knowledge processing framework, intends to close this gap and provides robots with the required information to perform their tasks. It builds on top of knowledge representation research, making the necessary adaptations to fit the robotics domain where typically much more detailed action models are needed. The core idea behind KNOWROB is to automatically adjust the execution of a robotic system to a particular situation from generic meta action models. The platform is validated with real robots acting in a kitchen environment with a strong focus on manipulation and perception. Beetz et al. combine KNOWROB with the usage of CRAM [46], which serves as a flexible description language for manipulation activities. CRAM in turn is used with the Semantic Description Language (SDL) which links capabilities with abstract hardware and software requirements through an

ontological model. As a result, symbolic expressions in CRAM can be grounded depending on the available hardware. CRAM is, however, not a planning system that can be used to solve arbitrary problems. Instead it can formulate a plan template for an already solved planning problem.

Meanwhile, reasoning in Q-ROCK aims at using planning, in particular Hierarchical Task Networks (HTNs), to generically formalize a problem in the robotic domain and to generate an action recipe as solution. HTN planning is an established technology with a number of available planners such as CHIMP [47], PANDA [48] or SHOP2 [49], but there is still no de-facto standard language comparable to the Planning Domain Definition Language (PDDL) [50] in the classical planning domain. Höller et al. [51] suggest an extension to PDDL hierarchical planning problems named Hierarchical Domain Definition Language (HDDL) to address this issue. Nevertheless formulating an integrated planning problem which includes semantic information remains an open challenge.

Approach & formalization

Top-Down: Identification of capable systems We start by describing the process of structural reasoning from entry point E2 into the Q-ROCK development cycle (see Fig. 1). The workflow for the top-down reasoning is depicted in Fig. 7.

To enter the cycle at E2 a user has to provide a description of an application problem to solve, i.e., defining tasks

that should performed and the application environment including the initial state. The problem is described with a general language and is firstly hardware agnostic. This means, neither does the application description explicitly state the use of a particular robot nor a robot type. While an input using natural language would be desirable for users to describe their application, Q-ROCK uses a planning language like PDDL or as directly machine readable format. Formulating the application problem firstly generically and secondly as hierarchical planning problem allows the decomposition into a sequence of atomic / primitive tasks, where $p \in P$ denotes a primitive planning task and P denotes the set of all primitive tasks. Q-ROCK extends state-of-the-art planning approaches by (a) introducing a semantic annotation for each primitive task, and (b) representing the domain description, i.e., all tasks and decomposition methods, with an ontology.

The semantic annotation of a primitive task comprises a constrained-based description of what a task does in the classical sense of planning effects, i.e., what it requires to start the execution as preconditions and the condition that have to prevail during an execution. All conditions including pre/prevail and post can be tested upon using a predefined set of predicate symbols, which describe the partial world state including environment state s^{obs} and robot state s . Hence, the semantic annotation of a primitive task also includes pre and prevail conditions that link to the state of hardware and software components.

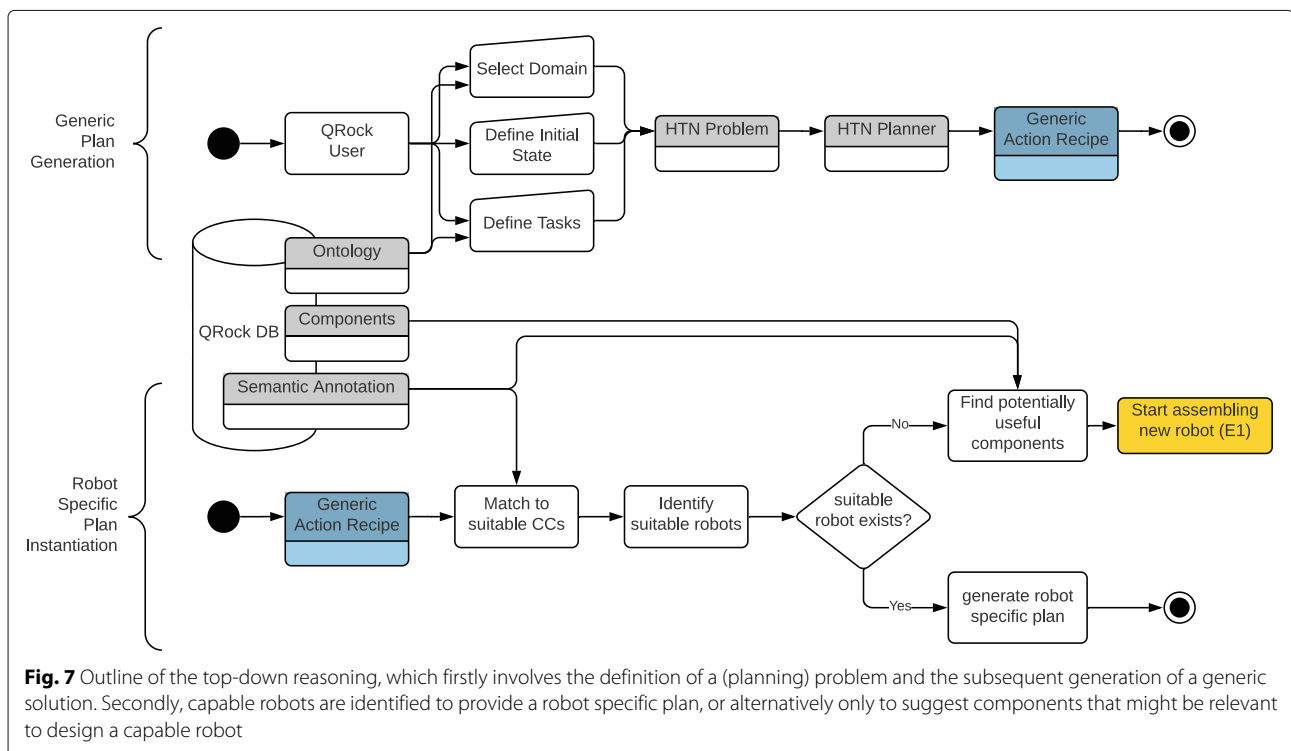


Fig. 7 Outline of the top-down reasoning, which firstly involves the definition of a (planning) problem and the subsequent generation of a generic solution. Secondly, capable robots are identified to provide a robot specific plan, or alternatively only to suggest components that might be relevant to design a capable robot

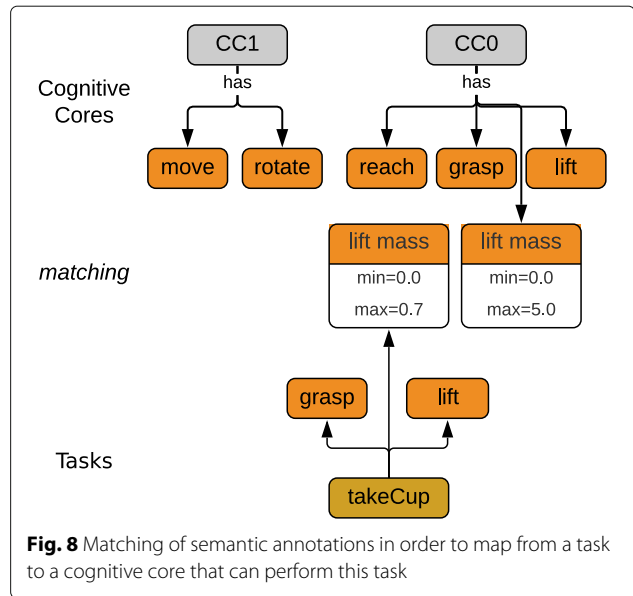
Definition (Semantically Annotated Primitive Task). A *semantically annotated primitive task* $p^+ \in P^+$ is a tuple of a primitive planning task p and a semantic annotation SA , so that $p^+ = (p, SA)$. P^+ denotes the set of all *semantically annotated primitive tasks*.

The top-down reasoning is based on a predefined planning vocabulary $\mathcal{V}_p = (P, C, \mathbf{d}, sa)$ to specify problems, here representing a particular planning domain description, where the vocabulary consists of primitive (P) and compound tasks (C), decomposition methods \mathbf{d} for compound tasks, and a mapping function $sa : P \rightarrow SA$, SA denoting the set of all semantic annotations. The top-down reasoning process is initially limited with respect to the expressiveness of this application specification language.

Transforming the user’s problem into a planning problem and solving it results in a collection of plans, where each plan in this collection represents a robot type agnostic solution. This does not imply, however, that the requested task is solvable with current available hardware. Each semantically annotated primitive task that is part of a solution has requirements for its execution including, but not limited to environmental, temporal, and hardware and software constraints. Therefore, an additional validation of these constraints has to be performed.

Requirements to execute a plan can be extracted from the semantic annotations belonging to all of its semantically annotated primitive tasks, in the simplest case by the use of labels which classify. Semantic annotations also describe cognitive cores as explained in “[Classification and annotation](#)” section. Such a description might be incomplete in the sense that it does not catch every detail of the behavior of a cognitive core, but it serves to outline the semantics in an abstract and also machine processable way. Furthermore, it allows to match semantic annotations of tasks against semantic annotations of known cognitive cores. Thereby identifying cognitive cores that can be used to tackle the stated problem (see Fig. 8). Each cognitive core maps to a single robotic system, but primitive tasks can map to different cognitive cores. Finally, a solution is only valid if a single suitable system which is capable to perform all tasks can be identified. While this concept of matching tasks and cognitive cores can also be used to map to multiple systems that cooperate to solve the stated problem, Q-ROCK focuses on single robots for now.

As outlined before, Q-ROCK aims at a planning approach which does not focus on a particular robotic systems, but provides abstract solutions. Although no specific robot types are considered, solutions still can comprise hardware requirements to solve a particular task. For instance, a requirement could be the emptiness of a gripper before starting a gripping activity. This particu-



lar precondition, however, implies also the availability of a gripper and thus restricts the applicable robot types that can be used to perform for this task to those that have a gripper. A selected target object might induce additional constraints for lifting mass or handling soft objects, so that only a particular type of gripper can be used.

Effectively, the following structural requirements exist for hardware and software components: 1. existence of hardware and software components in the system, and 2. particular (sub)structures formed by hardware and/or software components. Additionally, functional requirements exist which might imply structural requirements, so that functional requirements can be considered as higher-order predicates for tasks. These could be implemented similar to using semantic attachments for planning actions as suggested by [52]. Workspace dimensions and required maximum reach are examples of an extended task description, which limits the range of systems applicable for this task.

To create semantic annotations, Q-ROCK uses a corresponding language \mathcal{L} . Meanwhile, Q-ROCK uses an ontology to represent the vocabulary $\mathcal{V} \supset \mathcal{V}_p \cup \mathcal{V}_{SA}$ of this language, which combines the planning vocabulary \mathcal{V}_p and the semantic annotation vocabulary \mathcal{V}_{SA} which permits to specify components, labels (corresponding to behavior types) and constraints. While labels allow to classify and categorize behaviors, constraints allow to detail or rather narrow these behaviors further, e.g., manipulation with a constraint to manipulate a minimum of 100 kg mass poses a significant hardware constraint.

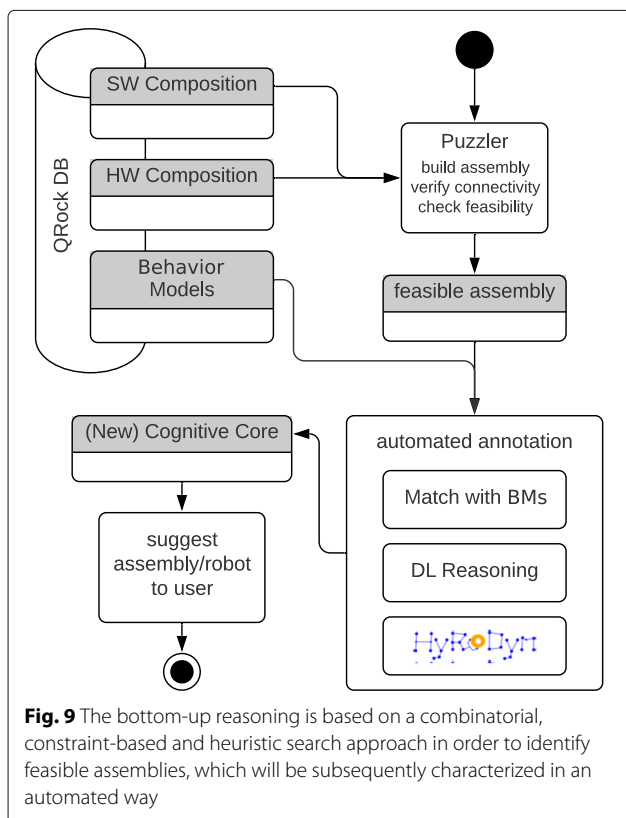
Semantic annotations characterize primitive tasks as well as behavior models and cognitive cores and are therefore essential to link Q-ROCK’s clustering step with top-down reasoning.

Bottom-Up: From Structure to Function While the top-down reasoning process tries to find suitable hardware for a given task, the bottom-up approach aims at finding the functionality or rather tasks that a component composition can perform. The bottom-up identification of a robot's function is based on a formalization introduced by Roehr [53] who establishes a so-called organization model to map between a composite system's functionality and the structure of components. Functionalities can be decomposed into their requirements on structural system elements. As a result, the known structural requirements for a functionality can be matched against (partial) structures of a composite system to test whether this functionality is supported. Figure 9 depicts the bottom-up reasoning workflow, where an essential element of the bottom-up reasoning is the identification of feasible composite systems or rather assemblies. The combination of components requires knowledge about the interfaces of the components and permits a connection between any two components only if their interfaces are compatible. Multiple interfaces might be available for connection and physical as well as virtual (software) interfaces can be considered. Roehr [53] limits the bottom-up reasoning to a graph-theoretic approach while excluding restrictions arising from the actual physical properties of the overall component, e.g., shape or mass. Q-ROCK will

remove that restriction and analyse the actual physical combinations of components as part of the so-called *puzzler* component. The *puzzler* component composes new assemblies from a known set of atomic components by creating links between compatible interfaces of atomic components. By using the existing D-Rock tooling, and extending the component model specification with ontological knowledge from Knowledge-based Open Robot vocabulary as Utility Toolkit (Korcut) [13], the newly defined assembly is loaded into blender and exported to Universal Robot Description Format (URDF) along with additional other material and sensor information. Subsequently, the new assembly can be physically validated and explored in simulation. Based on the URDF representation we will use Hybrid Robot Dynamics (HyRoDyn) [54] for characterizing the robot, e.g., by computing basic properties of the robot in zero configuration, and analysing configuration space, workspace and forces. Furthermore, the assemblies will be semantically annotated by (a) matching the structure to existing behavior models and (b) reasoning on the ontological description. An initial manual and later automated analysis of component structures in existing robotic systems can serve as a basis to identify generic design patterns in robotics systems. This can be used as heuristic to boost the bottom-up reasoning process. The bottom-up reasoning process is triggered from new additions of software and hardware to the database. Thereby, augmenting the Q-ROCK database by adding new cognitive cores helps to increase the options for solving user problems.

Results: a use case scenario

The Q-ROCK development cycle combines existing AI technologies in order to (a) simplify the robot development process and (b) exploit the full capabilities of robotic hardware. In this section we provide a qualitative analysis of the Q-ROCK development cycle on the basis of three complementary use cases which allow us to illustrate the concept and locate conceptual as well as practical issues. The use cases are: (1) system assembly, and then using the assembly for (b) solving a task and (c) solving a mission. For the system assembly the workflow starts at entry point E1 (see Fig. 1), namely constructing a robot model, exploring its movement capabilities, clustering these capabilities and generating a cognitive core. The framework is designed to handle all kinds of capabilities that can be described as trajectories in the combined robot and world state. For this example and initial evaluation we focus, however, on movement capabilities. We use three test systems - a 3-DOF robotic arm, a wheeled mobile base, and a combination of both - and create a cognitive core *reach*, which moves the end effector from a given start configuration to a chosen target position in task space.



The schematic workflow is visualized in Fig. 10, while Fig. 11 highlights steps of this workflow which are triggered through our implemented website. Several preparatory steps are, however, assumed and required for the workflow to run including the definition of (a) feature spaces and (b) behavior models. We validate the creation of the cognitive core and integration of the system, by retrieving the cognitive cores again via its labels from the database.

Preparation

We predefine feature spaces which are used in the clustering step and for defining a behavior model. Note, that this is a simplification which we will address in future implementation, e.g., by the development of automated feature learning approaches.

Defining Feature Spaces To characterize a *reach* movement the feature spaces firstly permit to extract the starting state and the end state of trajectories. Further qualities of a trajectory such as its directness are also included to penalize deviations of a trajectory from the direct route. As feature spaces we use:

1. F_{start} with label 'start state' and the transformation function $f_{start} : L \rightarrow S$ which maps a given trajectory $l^T \in L$ to its start state $s^0 \in S$,
2. F_{end} with label 'end effector end state' and the transformation function $f_{end} : E \rightarrow S$, which maps a given end effector state trajectory to its final state e^T ,

3. F_{dir} with label 'end effector directness' and the transformation function $f_{dir} : L \rightarrow \mathbb{R}$ and

$$f_{dir} = \frac{|s_e^0 - s_e^T|}{\sum_{i=0}^{T-1} |s_e^{i+1} - s_e^i|} \in (0, 1]$$

Note that the features F_{end} and F_{dir} use the end effector position, which we assume is part of the world state S^{obs} , whereas F_{start} operates on the internal state S^{rob} of the robot.

Defining Behavior Models A behavior model provides the high level abstraction for a behavior, thereby collecting the essential characteristics. In the case of the *reach* behavior the key characteristic is the directness, which we expect to be high so that it is bound to a minimum and maximum degree. Meanwhile the start and the end state are variable, since the *reach* behavior needs to be applicable in a range of situations with different target poses. Hence, start and end state can be viewed as general input parameters to the behavior model.

label : reach
 constraints : F_{dir} : min : 0.8 max : 1.0
 F_{start} : variable
 F_{end} : variable

All defined feature spaces, the behavior model and the semantic annotation of the behavior model, which contains labels and constraints, are stored in the database to be accessible for all development steps.

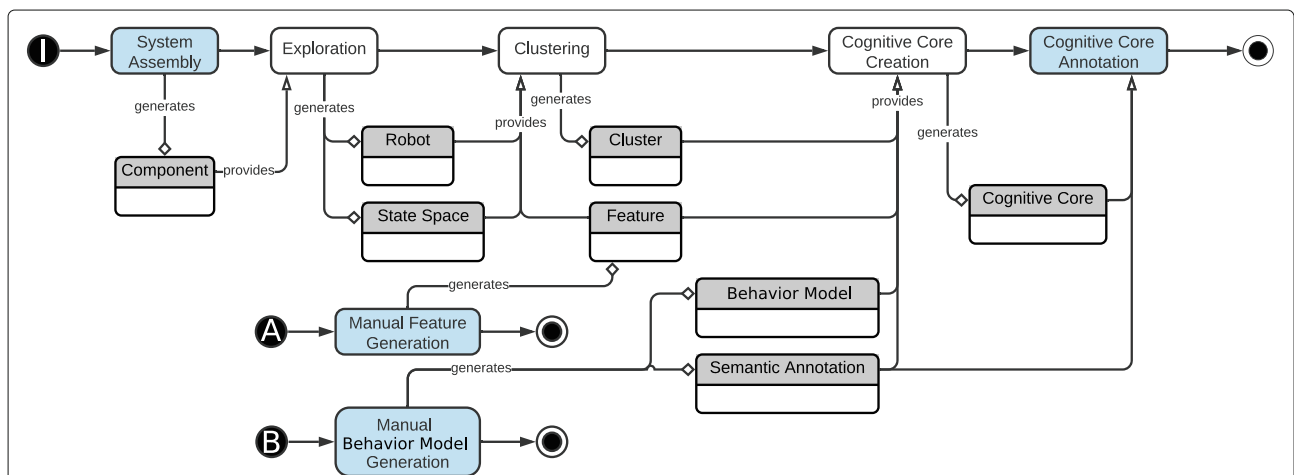


Fig. 10 Workflow for our exemplary use case. Steps with required user interaction are shaded blue. Relations between entities are avoided for clarity. The full entity relation diagram is visualized in Fig. 2. We start with entry point A to manually define features. Then proceed with entry point B and manually define a behavior model. Then the actual work flow starts at I. A system is assembled using components from the database, generating a new component. This new component is passed to the exploration step, which generates a robot and a state space entity in the database. Part of the robot entity is the capability function, which is used in the clustering step, along with information about the state space of the robot and the features in which to cluster. The clusters that are generated are used in cognitive core creation, which grounds the previously defined behavior model for this robot. In the cognitive core annotation step, the semantic annotation inherited from the behavior model is reviewed and possibly extended with robot-specific information by a human observer

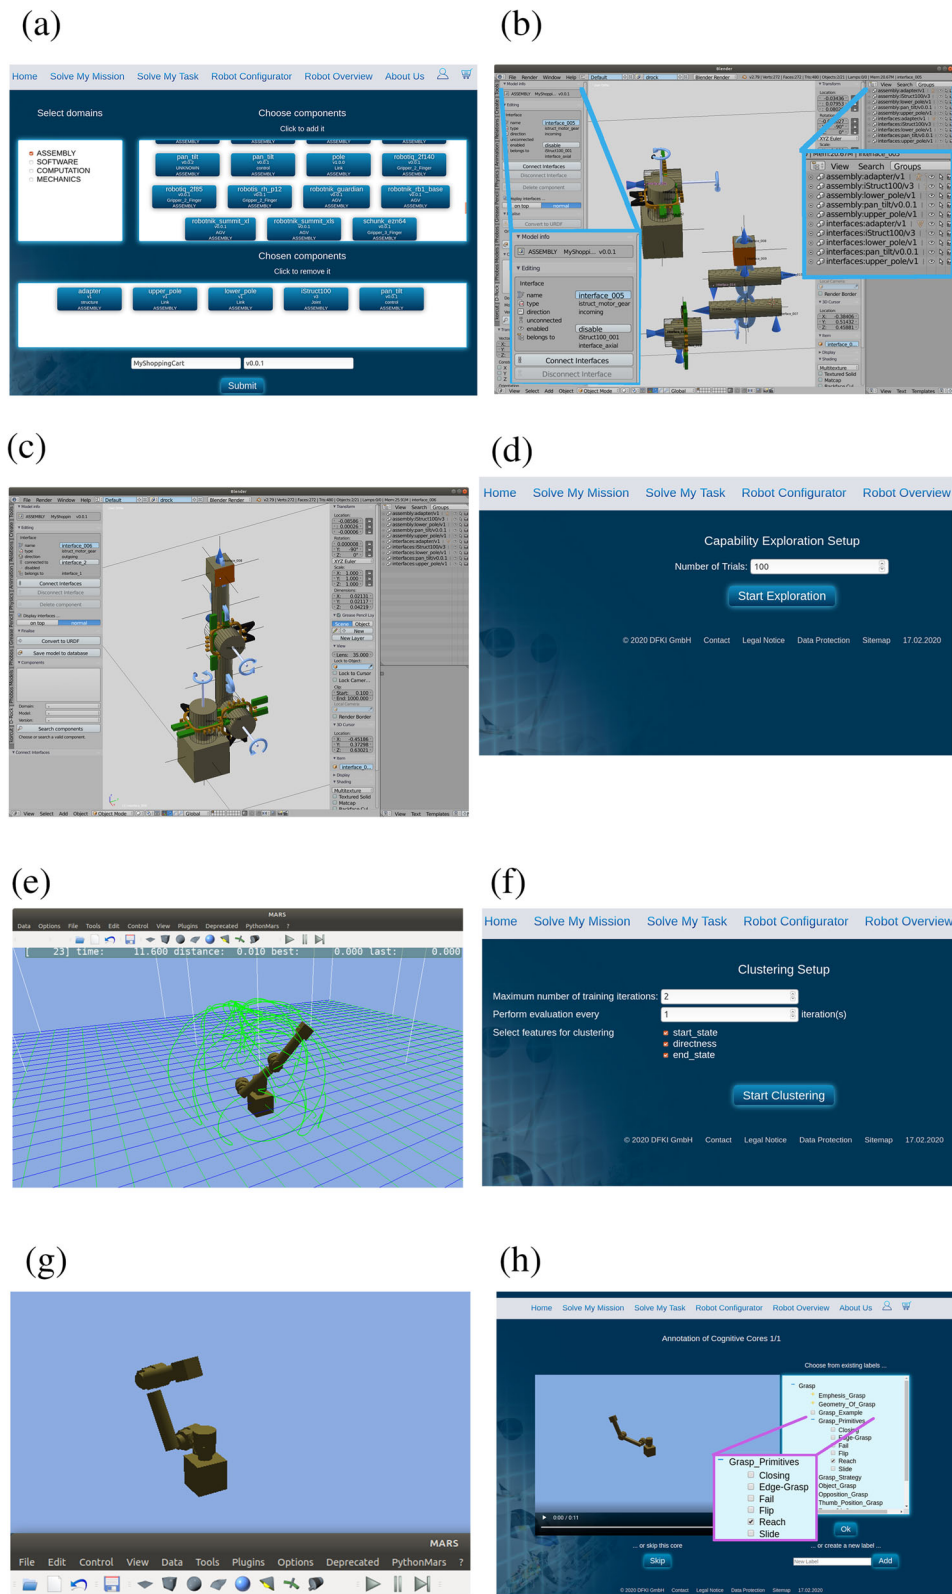


Fig. 11 The *Robot Configurator* workflow takes advantage of exploration and clustering and allows to construct a robot first, which will then automatically be explored and annotated. **a**) Choosing the components **b**) Assembling components with the help of Blender and Phobos **c**) Saving the final assembly into the database **d**) Parametrize the exploration **e**) Explore the capability of the new assembly **f**) Clustering is applied and cognitive cores identified from the existing behavior model labelled reach **g**) A video of the identified cluster performance is automatically rendered **h**) A user can watch videos of a cognitive core and an- notate

Note that defining the behavior model is an essential, but currently also a limiting requirement, since the exploration of behaviors can only cover these predefined models.

System assembly

To start the Q-ROCK development cycle at entry point E1 a robotic system is required. Predefined robots or rather existing assemblies can be used to start the exploration. One of the major motivations of the Q-ROCK development cycle is, however, the capability to explore any kind of hardware designs / assemblies and thereby support an open robot design process.

The so-called *Robot Configurator* workflow permits a user to create a robotic system in a simplified way, by combining a set of components that are defined in the database. Figure 11 illustrates the steps. Firstly, a user selects the desired items which are needed to build the robot and puts them in a shopping basket (see Fig. 11a). For the *reach* example an assembly is built from the following items: 1. pan tilt unit, 2. lower pole, 3. joint motor, 4. upper pole, and an 5. end effector. After the selection has been completed, a CAD editor is started with the selected items being already loaded (see Fig. 11b). We use the open source editor Blender² in combination with the extended functionality of the Phobos plugin [55] and another custom plugin to interface with the database. The user can build the desired system by selecting interfaces in the GUI and request to connect components through these interfaces, which is only possible if the selected interfaces are compatible. Component interfaces and their compatibility are defined in hand curated ontologies. The overall procedure requires only very limited editing competencies of a user, thus significantly lowering the entry barrier for physically designing a robot. Once the final system has been assembled, the user can save the new design to the database (see Fig. 11c).

Exploration For the exploration of movement capabilities we chose a capability function model where the parameters θ correspond to polynomial parameters. The parameters define an intended joint trajectory for all joints of the robot. For a single joint the trajectory is defined by:

$$q(\phi) = -\theta_0(\phi - 1) + \theta_1\phi + \sum_{i=2}^{N_\theta-1} \theta_i (\phi^{i-1} - 1) \phi$$

where $\phi = \frac{t}{T}$ is a phase and T the length of the trajectory. The first parameter θ_0 corresponds to the start position and θ_1 to the final position. The number of parameters per joint was set to $N_\theta = 5$.

The Q-ROCK development cycle is also compatible with other motion representations such as splines and dynamic movement primitives. Meanwhile, this polynomial representation has the advantage that higher order parameters only contribute if their value is different from zero. This potentially eases the clustering process on this parameter space and it makes combining parameter sets with different N_θ straightforward.

The intended joint trajectories are passed to a controller which returns the robot specific actions that are necessary to follow the desired trajectory as closely as possible. The controller, where the desired states are specified by a trajectory, takes the role of the capability function *cap* in this context. The actions applied to the real system will finally result in a capability.

This capability function model ensures that the mapping from Θ to the capabilities is locally smooth, i.e., small changes of the parameters will lead to small changes in the corresponding trajectory. This is an important property for modelling clusters of trajectories in the Classification and Annotation process (“[Classification and annotation](#)” section).

To test the exploration approach, as described in “[Exploration](#)” section, 10^6 trajectories were generated for a robotic 3-DOF arm. Each trajectory has a length of $T = 4$ s. Five parameters specify the motion for every degree of freedom, i.e., each parameter lies within $[-\pi, \pi]$, resulting in 15 parameters for the whole 3-DOF robot. The parameter space has been sampled uniformly. The same procedure was applied to a mobile base (see Fig. 12 right) and a combined system of mobile base and 3-DOF arm (see Fig. 13 right). For the mobile base, the polynomial trajectories were applied to the velocity of the four wheels with a parameter range of $[-2\pi, 2\pi]$. Three parameters were used to parametrize the motion, leading to 12 parameters in total. The trajectory length was set to $T = 20$ s. The combined system thus has a total of 27 parameters. The validation model has been implemented as a fully connected neural network with four hidden layers and 22,102 parameters in total. Training for the 3-DOF arm was performed for 20 training steps with a batch size of 100. After training, the validation model had an accuracy of $96.6 \pm 0.4\%$ on unseen data of 10^4 trajectories for predicting whether a parameter set corresponds to a valid motion.

Clustering For clustering, the same capability set as for training the validation model was used. Without loss of generality, we implemented standard k-means clustering as the clustering strategy [56], choosing 50 clusters in F_{start} and F_{end} and 5 clusters in F_{dir} . The cluster numbers were chosen manually to assure a large enough size of each cluster for generative model training, while also providing a sufficient resolution in the respective feature

²<http://www.blender.org>


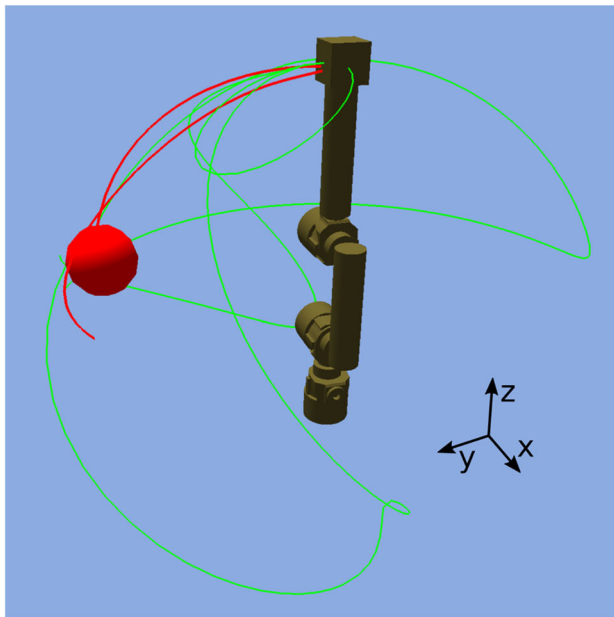


Fig. 12 Reach cognitive core for two robotic systems. Left: End effector trajectories generated by sampling from the reach core of a 3-DOF arm (red lines) are shown with the requested target area of the behavior (red ball) at $(x, y, z) = (0.1, 0.3, 0.2)$ m. For comparison, samples from the unconstrained reach cognitive core - without the constraint on directness - are plotted in green. Right: Same as left, but for a wheeled mobile base, with a target location at $(x, y, z) = (-1, 1, 0)$ m




Fig. 13 Reach cognitive core samples for a combined system of 3-DOF arm and mobile base. The red lines visualize the end effector motion of the arm of the combined system, thus including the motion of the mobile base. The target location is marked by the red ball at $(x, y, z) = (1.3, -0.1, 0.5)$ m, which is out of range of the arm alone, would the mobile base not move as well. The arm has a length of ≈ 0.4 m. Note that the directness constraint only applies to the end effector motion, which is the end effector of the arm for this combined system, and not the task space motion of the base. For comparison, samples from an unconstrained reach cognitive core are plotted in green. As expected, samples from this unconstrained core reach the target location more indirectly

space. As generative models, we implemented neural ordinary differential equation (ODE) based normalizing flows [57, 58]. We made this choice over alternative, more classical methods for density estimation, such as Gaussian Mixture Models or variational autoencoders, since we found that the parameter distributions can be too complex to be reasonably captured by GMMs, and that the normalizing flow models were very robust over a big range of hyperparameters. We used a batch size of 128, 1000 training iterations, and a network layout of two fully connected hidden layers with 64 neurons each. After training, the model accuracies are $91 \pm 2.1\%$ for F_{start} , $87 \pm 3.3\%$ for F_{dir} , and $59 \pm 4.1\%$ for F_{end} for the 3 DOF arm, $89 \pm 3.2\%$ for F_{start} , $73 \pm 4.3\%$ for F_{dir} , and $47 \pm 3.6\%$ for F_{end} for the mobile base, and $82 \pm 4.2\%$ for F_{start} , $55 \pm 3.9\%$ for F_{dir} , and $32 \pm 5.1\%$ for F_{end} for the combined system. The model accuracies were determined by sampling parameters from the generative models, simulating the corresponding trajectories, and calculating the feature values. The accuracy is the percentage of samples that are assigned to the cluster the model was trained on. Errors were calculated from 10 training runs per model with randomized initial weights. Cluster entities for the robot in this example are generated and stored in the database for the identification of cognitive cores.

Cognitive Core Creation Based on the defined behavior models, the cognitive core is instantiated for all three robots after exploration and clustering, and inherits the

label *reach* (see also Fig. 6). For all systems, one cluster in the constrained feature space F_{dir} was found with a centroid value of ≈ 0.9 . Since the cluster fulfils the behavior model's minimum and maximum value constraints for this feature space, the cluster is linked to the cognitive core.

Cognitive Core Sampling One important aspect of the cognitive core is the ability to sample trajectories that conform to the underlying behavior model. In the current example, particle swarm optimization algorithms are used for the global optimization of the three feature models that are combined in the *reach* cognitive core, while constraints and feature inputs are weighted equally. Due to the stochastic nature of this process, several sub-optimal solutions are generated during sampling that represent the corresponding behavior.

To evaluate the specificity of the clustering procedure, we also define an alternative *reach* cognitive core without the directness constraint. This cognitive core will generate behavior that reaches an end effector end position in any way possible. Figure 12 shows samples from the constrained and the unconstrained *reach* cognitive core for both the 3-DOF arm and the mobile base. The trajectories were obtained by simulating the parameters sampled from the cognitive cores. All simulations and visualizations were performed in the MARS simulator [59], which is integrated in the Q-ROCK workflow. As expected, the constrained *reach* cognitive core generates direct motion to the desired goal, whereas samples from the unconstrained one are more indirect. Note, however, that these indirect trajectories can be useful in different settings, e.g., when parts of the task space are obstructed. This example also illustrates that samples from the cognitive cores do not exactly match the desired start and end position, and also show variability along the trajectory. Whereas this might seem counter-intuitive from a pure controlling standpoint, cognitive core sampling is deliberately probabilistic to allow generation of all possible behaviors that match the corresponding behavior model description, and not one single optimal trajectory.

To demonstrate the applicability of our approach to more complex systems, we show samples from the *reach* cognitive core of a combined system of 3-DOF arm and mobile base in Fig. 13. Here, combinations of motion of the mobile base and the attached arm are required to reach the target location of the end effector. As in the previous examples, the *reach* cognitive core successfully generates direct motion towards the target location.

Another way of generating more complex behavior for a system of already explored components is to use the previously generated cognitive cores for each subsystem. Figure 14 shows the parallel execution of both *reach* cognitive cores on each subsystem. The cognitive cores of the arm and the mobile base are sampled individually to

reach a target with the end effector of the arm that would lie outside of a fixed arm's task space. As expected, the assumption that the task space trajectory of the combined system can be simply composed from the trajectories of the subsystems does not generally hold. However, the sampled motion can be a good first approximation of the actual behavior when the dynamical coupling between subsystems is not too strong.

Cognitive Core Annotation For the annotation step, the cognitive core is executed several times with different variable inputs to F_{start} and F_{end} . Videos of the performance of the cognitive cores are generated and shown to a user, who can confirm the selected labels for the cognitive cores or (re)assign labels. For this demonstration the user approves and sticks to the label *reach* for the cognitive core, which has been inherited from its behavior model.

Solving a task

Cognitive cores are semantically annotated in order to provide a high-level description or rather specification of their performance. A user is not necessarily interested in designing new systems, but will typically first search for available robots which can solve the task at hand, or show similar performances. For this use case, we designed the workflow named *Solve My Task* (see Fig. 15), where a user selects a combination of semantic labels from the existing ontology and matches them against existing cognitive cores in the database. Before the user has to make a final choice, the performance of each identified cognitive core can be inspected through the previously rendered videos. Here, the explored *reach* cognitive core can successfully be retrieved and visualized to the user.

Solving a mission

The final use case, illustrated in Fig. 16, deals with solving a user's application scenario, in the following referred to as *mission*. A mission can range from a single robot action to a complex plan involving multiple actions that need to be sequenced. A mission with sequential actions can be composed through our web interface, based on a set of predefined - yet generic - actions: grasp, navigate, perceive, pick, reach, release. For this evaluation we select the *reach* action, which maps to a requirement for cognitive cores with a semantic annotation including the label "reach", so that as an intermediate result the previously identified core can be picked. The cognitive core is linked to the design of the 'NewShoppingCart', which can now be considered a suitable robot system to perform the mission. Therefore, this custom design is the final suggestion of the Q-ROCK development cycle to solving this mission.

Summary

The complementary use cases show the working of key elements of the Q-ROCK development cycle. Assemblies




Fig. 14 Parallel execution of *reach* cognitive cores of subsystems. The cognitive cores were trained on each system individually and are samples of the same cognitive cores shown in Fig. 12. Left: The red lines show the end effector motion of each system in isolation, the magenta line the expected motion of the arm end effector of the combined system under the assumption that each subsystem behaves as in isolation. The red ball marks the requested target for the combined motion at $(x, y, z) = (1, 2, 0.5)$ m, outside of the task space of the arm alone, which has a length of ≈ 0.4 m. The actual motion of the base and the arm of the combined system is shown by the blue lines. Right: Same as left, but for a different target position $(x, y, z) = (1.2, 1.7, 0.5)$ m and different samples from the cognitive cores for each subsystem, in which the dynamics of the subsystems lead to a non-negligible interaction

can be composed in a simple manner by reusing predefined components from the Q-ROCK database. The assemblies are explored in simulation to identify behavior and subsequent clustering of trajectories allows to extract cognitive cores. Semantic annotations bridge the gap between cognitive cores and planning tasks, so that the loop between the bottom-up system analysis and top-down requirement-based robot design is closed. Overall, although the use case for the *reach* behavior is simple, this evaluation shows a fully integrated workflow implementation. It validates the feasibility and working of the concept as a first step towards automatically developing complex robot behavior.

Discussion

In this paper, we introduced the formal concepts behind Q-ROCK' and presented use cases to demonstrate our approach to solving problems in a combination of bottom-up and top-down solving. The use cases show that the integrative workflow has been established, although several challenges remain. In the following, we will discuss key aspects of the workflow in more detail.

Exploration

The exploration framework has been successfully implemented and allows the automated data generation

for assemblies built with the robot configurator. The implementation is designed in a modular way, so that minimal effort is needed to switch between capability function models.

In the current state the exploration is done in the configuration space of an assembled system. This approach is not scaling well if the system complexity is increasing, such as for systems of systems. Whereas for the combined system illustrated in Fig. 14, the full exploration approach is still feasible, it will break down at some complexity level. To allow an application of the exploration approach to these complex cases we have to make use of the knowledge already generated for the subsystems. For the simplest concept to reuse the subsystem capabilities, one would sample capabilities out of the subsystems clusters to generate capabilities for the combined system and thus form an exploration dataset. The concept development and evaluation is an ongoing process.

After the training process, an implemented validation model can be queried for the information whether a capability can be executed by the robot. However, selecting a capability and then checking it is inconvenient for the clustering process. For this reason, in the next step we are looking into inverting this validity information, i.e., mapping all valid (and only the valid) capabilities into a

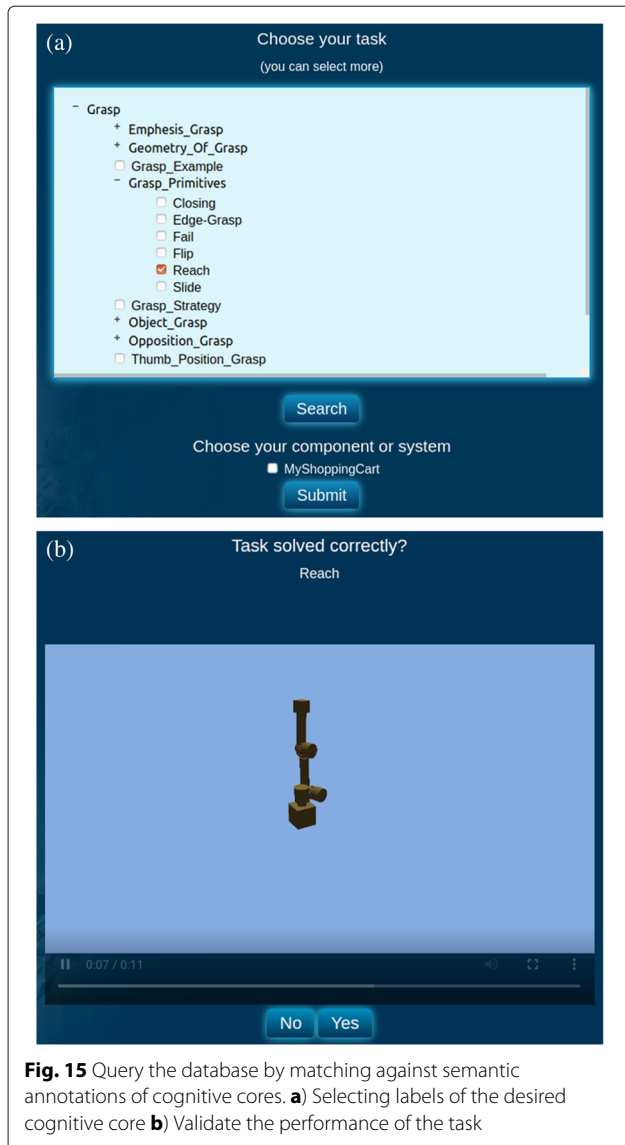


Fig. 15 Query the database by matching against semantic annotations of cognitive cores. **a)** Selecting labels of the desired cognitive core **b)** Validate the performance of the task

continuous parameter space, on which the clustering can operate.

The most challenging problem for the exploration is to make these approaches scale with increasingly complex systems. In order to deal with this challenge we plan to apply more sophisticated search strategies. As the exploration is supposed to be task independent we intend to use intrinsic motivations [22] to explore the search space in a structured manner.

Another challenge is the exploration of perception capabilities which is considered in the theoretical framework, but requires non-trivial environments to perform the exploration. Designing test environments to allow a mostly task independent exploration is one main challenge, besides the fact that for perception the capability space dimensionality is even higher compared to kinematic and dynamic exploration.

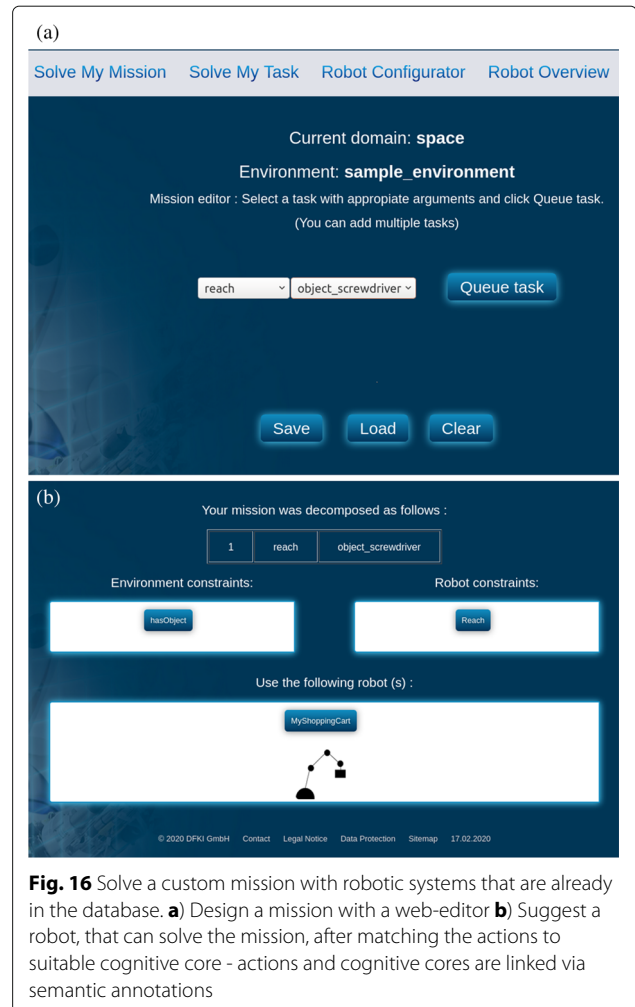


Fig. 16 Solve a custom mission with robotic systems that are already in the database. **a)** Design a mission with a web-editor **b)** Suggest a robot, that can solve the mission, after matching the actions to suitable cognitive core - actions and cognitive cores are linked via semantic annotations

In parallel to the exploration approach, an introspection into failure cases is envisaged via a hierarchical capability checking framework which can (a) detect whether a given action is feasible on the robot and (b) pin-point the reasons of infeasibility. This problem is especially interesting for mechanisms with closed loops, e.g., parallel robots or serial-parallel hybrid robots [60]. We plan to exploit knowledge about the kinematic structure of the robot, its various physical properties, and analytical mappings between different spaces (actuator coordinates, generalized coordinates and full configuration space coordinates) by using HyRoDyn which is under active development in Q-ROCK.

Classification and annotation

Following exploration, capability clustering and the application of cognitive core and behavior model formalizations have also been successfully implemented, while revealing interesting challenges.

Clustering is an important step in the Classification and Annotation workflow, since it increases the granularity

of the latent space on which the generative models are trained. Whereas the current implementation, using k-means clustering, has the advantage of being robust and well developed, a shortcoming is that if more finely spaced clusters are necessary for a specific behavior model, corresponding regions of the feature space have to be clustered again and cluster models trained anew. We are thus pursuing the integration of more sophisticated generative model learning approaches into the workflow that retain more accessible information about the underlying data distributions, such as arbitrary conditional flows [61].

Behavior models are also an important aspect in this part of Q-ROCK. Whereas only hand-crafted behavior models and feature spaces have been tested to date, we aim at a more automated approach in the future. We actively research the application of variational autoencoders to world state trajectory data, and how well features found in this way can be semantically interpreted. Furthermore, we are working on automated extraction of behavior model definitions, both from human demonstration data and from modelling human evaluation functions.

Although our goal is to increase the level of automation in the future, we still see human labelling as a crucial backstop in the cycle to give meaning to the explored data and to introduce steps for revision as part of the bottom-up path. Recent work has also shown that pure automation based AI approaches can be inferior to an interactive human in the loop in complex reasoning tasks [62]. We thus see leveraging human semantic knowledge rather to be a feature of our approach than a shortcoming.

An interesting theoretical problem is defining where the domain of cognitive cores ends, and the domain of planning begins, i.e., up to which behavioral complexity level cognitive cores can be reasonably defined. The cognitive core formalism is purposefully flexible enough such that planning algorithms can be expressed, thus there is no clear limitation imposed on the formal side. A related challenge is parallel execution of cognitive cores of subsystems, such as demonstrated in Fig. 14. Whether learning based techniques that are warm-started with initial guesses from individual cognitive core samples, higher level policy training on the established latent space of simple policies similar to [32], or more reasoning oriented approaches will prove more effective in our workflow is not finally resolved. An alternative solution is only using cognitive cores of a full system and using its subsystems' clusters for the exploration of the full system, as discussed in section "Exploration".

Reasoning

To effectively exploit explored cognitive cores and behaviors, we suggest a high-level planning based approach in this paper. Since the existing decomposition of problems

is currently depending on the planner's domain definition, we also provide an interface to create new missions. The current use case does not challenge this interface, since only very simple tasks decompositions are required. For increasingly complex problem descriptions, this mission description interface needs to remain not only intuitive, but also sufficiently expressive. Furthermore, this interface needs to be extensible through a growing vocabulary. Another challenge remains for the application of a found solution and exploitation of existing behaviors to address a user's high-level problem. The framework provides first a mostly generic and robot-agnostic solution. This solution has then to be mapped to the finally selected robot, i.e., this robot needs to execute the solution. To achieve the latter, we have to run our cognitive core in robot-specific contexts, e.g., exploiting existing robotic frameworks such as ROS. Thus, a grounding of generic solutions to selected environments and newly designed robots has to be performed. Semantic annotations are key elements for this mapping since they provide the essential glue between reasoning and cognitive cores, where the vocabulary is defined by our ontology. We still need to enrich and revise the structure of this ontology based on human feedback processes of annotating cognitive cores. Hence, developing a sufficiently expressive semantic annotation language remains a further challenge for the reasoning part.

The size of the database or rather number of components and options to combine components in new systems leads to another, combinatorial challenge. Here, we need to find effective heuristics in the context of the puzzler development.

Outlook

From the view point of Q-ROCK as a whole, we see two major challenges arising for future work. On the one hand, the system requires a rich database of annotated components, i.e., single parts or already simple robots, together with a design flow to create new assemblies, in order to generate a significant added value for users. A proposal for modelling components in such a design flow has been developed in the predecessor project D-ROCK and is described in section "Modelling robot composition". With this, Q-ROCK has to be thoroughly tested, such that new robotic devices are created and many cognitive cores are built, which in turn fosters an enriched ontology to also interact with the user. During this process, it is important to evaluate the results of Q-ROCK in terms of completeness of found behaviors, as well as stability and robustness of the underlying representations.

A related challenge is the introduction of Q-ROCK to a considerable number of users to start forming a community. As a first step towards this goal, we intend to publish the parts of Q-ROCK open source, and a

strategy for addressing the robotics community is currently being formulated. The more users interact with Q-ROCK, and thereby also enrich the database, the more individual users will benefit and the more versatile it will become.

Taking a further step back, the Q-ROCK system is part of a greater development cycle in the X-ROCK project series. In D-Rock, the groundwork was laid for simplified modelling of robot parts and construction of robots based on well defined interfaces. In Q-ROCK, robots are enabled to explore possible behaviors. In future projects, beyond the scope of currently ongoing research, we plan to tackle questions regarding combinations of systems and their respective cognitive cores, behavioral interactions between humans and robots or groups of robots, and fine-tuning of behaviors for specific contexts.

Conclusions

The fundamental idea behind the Q-ROCK approach is to integrate and extend existing methods in AI, both on the symbolic and the sub-symbolic level, and to implement a framework that assists users to solve their intended task with an existing or novel robot. To achieve this, a central challenge is a unifying concept and theoretical foundation to (a) integrate all components in order to realize the Q-ROCK development cycle (given in Fig. 1), and (b), to have a clear definition of interfaces to extend the existing cycle or even replace single components with new approaches. Q-ROCK focuses on this integration to set up a new way of designing complex robots with the help of AI and with the knowledge that previous designers contributed to the knowledge base.

In this paper, we made an essential step by introducing the conceptual framework as a basis and integration platform for all subsequent work. Modularity is a key feature of our approach and allows embedding of alternative solutions for each stage. In the future, we expect competing or continuously improving implementations for each of the stages of the development cycle. With the use cases presented in this paper, we already demonstrated the functional coupling of all steps in Q-ROCK. In particular, the example shows that a model of annotated hardware can be used to successfully generate simple robotic capabilities (starting at E1 in the cycle in Fig. 1), which can be successfully clustered and annotated to generate a cognitive core. Hereby, a link is established from an exploration on the sub-symbolic level to a representation containing a semantic label, such that semantic input from a user can be made on which reasoning is performed. Q-ROCK is unique in this way: goal-agnostic capabilities are cast into a broader semantic framework, and sub-symbolic and symbolic levels in AI are integrated.

Abbreviations

AADL: Architecture analysis and design language; AWS: Amazon web services; BM: Behavior model; CAP: Capability; CC: Cognitive core; CF: Capability function; CFM: Capability function model; ER: Entity relationship; FF: Feature function; FS: Feature space; HDDL: Hierarchical domain definition language; HTN: Hierarchical task network; HyRoDyn: Hybrid robot dynamics; KR&R: Knowledge representation and reasoning; ODE: Ordinary differential equation; PDDL: Planning domain definition language; ROS: Robot operating system; SA: Semantic annotation; SDL: Semantic description language; UML: Unified modelling language; URDF: Universal robot description format

Acknowledgements

We thank all members of the Q-Rock development team and the internal reviewers for their valuable feedback on preliminary versions of this paper.

Authors' contributions

T.R. wrote the manuscript, developed software and theory. D.H. wrote the manuscript, developed software and theory. H.W. supervised the project, provided text for "Introduction" and "Discussion" sections. F.W. developed software, theory and provided text for "Exploration" and "Discussion" sections. M.S. developed software, theory and provided text for "Modelling robot composition" section, developed theory for 3. O.L. developed software, theory and provided text for "Reasoning" section. M.L. developed software and theory and provided text for "State of the art", "Exploration" and "Discussion" sections. S.K. provided text for "Discussion" section. S.S. supervised the project, provided text for "Introduction" and "Discussion" sections. F.K. conceived of the original idea, supervised the project, provided text for "Introduction" and "Discussion" sections. The author(s) read and approved the final manuscript.

Authors' information

(Optional - no information provided)

Funding

This research and development project is funded by the German Federal Ministry of Education and Research under grant agreement (FKZ 01IW18003). Open Access funding enabled and organized by Projekt DEAL.

Availability of data and materials

The datasets generated and analysed for the presented use case scenario are not publicly available.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

All authors gave their consent to publish this version of the manuscript.

Competing interests

The authors declare that they have no competing interests.

Author details

¹DFKI GmbH Robotics Innovation Center, Bremen, Germany. ²Institute for Communication Technology, Department of Information Technology, Dortmund University of Applied Sciences and Arts, Dortmund, Germany. ³AG Robotics, Department of Mathematics and Computer Science, University of Bremen, Bremen, Germany.

Received: 20 August 2020 Accepted: 30 April 2021

Published online: 05 July 2021

References

1. Stone P, Brooks R, Brynjolfsson E, Calo R, Etzioni O, Hager G, Hirschberg J, Kalyanakrishnan S, Kamar E, Kraus S, et al. Artificial intelligence and life in 2030: the one hundred year study on artificial intelligence. Technical report, Stanford University. 2016. https://ai100.stanford.edu/sites/g/files/sbybj9861/f/ai100report10032016fml_singles.pdf. Accessed 31 May 2021.
2. Siciliano B, Khatib O. Springer Handbook of Robotics. Berlin: Springer; 2016. <https://doi.org/10.1007/978-3-540-30301-5>.

3. Christensen HI. A roadmap for us robotics – from internet to robotics. 2020 Edition. Technical report, University of California San Diego, Computing Community Consortium, University of Massachusetts Lowell, University of Illinois, Urbana Champaign, University of Southern California. 2020. <https://www.therobotreport.com/wp-content/uploads/2020/09/roadmap-2020.pdf>. Accessed 31 May 2021.
4. Roehr TM, Harnack D, Lima O, Hendrik W, Kirchner F. Introducing Q-Rock : Towards the Automated Self-Exploration and Qualification of Robot Behaviors. In: ICRA Workshop on Robot Design and Customization. Montreal; 2019. Available at https://www.dfki.de/fileadmin/user_upload/import/10350_20190501_roehr_introducing_qrock.pdf.
5. Wiebe F, Kumar S, Harnack D, Langosz M, Wöhrle H, Kirchner F. Combinatorics of a discrete trajectory space for robot motion planning. In: 2nd IMA Conference on Mathematics of Robotics. Springer; 2021. accepted.
6. D-Rock. 2018. <https://robotik.dfki-bremen.de/en/research/projects/d-rock.html>. Accessed 31 May 2021.
7. Ha S, Coros S, Alspach A, Bern JM, Kim J, Yamane K. Computational design of robotic devices from high-level motion specifications. *IEEE Trans Robot.* 2018;34(5):1240–51.
8. Mansard N, DelPrete A, Geisert M, Tonneau S, Stasse O. Using a Memory of Motion to Efficiently Warm-Start a Nonlinear Predictive Controller. In: 2018 IEEE International Conference on Robotics and Automation (ICRA); 2018. p. 2986–2993. <https://doi.org/10.1109/ICRA.2018.8463154>.
9. Amazon Web Services (AWS). 2020. <https://aws.amazon.com>. Accessed 31 May 2021.
10. Neurobotics. 2020. <https://neurobotics.net/>. Accessed 31 May 2021.
11. Human Brain Project. 2020. <https://www.humanbrainproject.eu/>. Accessed 31 May 2021.
12. Tinkerbots. 2020. <https://www.tinkerbots.de/>. Accessed 31 May 2021.
13. Yüksel M, Korcut Ontology Family. <https://doi.org/10.5281/zenodo.4457562>. Accessed 31 May 2021.
14. Feiler PH, Lewis B, Vestal S, Colbert E. An Overview of the SAE Architecture Analysis & Design Language (AADL) Standard: A Basis for Model-Based Architecture-Driven Embedded Systems Engineering. In: Dissaux P, Filali-Amine M, Michel P, Vernadat F, editors. Architecture Description Languages. IFIP WCC TC2 2004. IFIP The International Federation for Information Processing, vol 176. Boston: Springer; 2005. https://doi.org/10.1007/0-387-24590-1_1.
15. Perrotin M, Conquet E, Delange J, Tsiodras T. Taste: An open-source tool-chain for embedded system and software development. *Embed Real Time Syst.* 2012. Available at <https://hal.archives-ouvertes.fr/hal-02191871/document>.
16. Scioni E, Huebel N, Blumenthal S, Shakhimardanov A, Klotzbuecher M, Garcia H, Bruyninckx H. Hierarchical hypergraphs for knowledge-centric robot systems: a composable structural metamodel and its domain specific language npc4. *J Softw Eng Robot.* 2016. https://doi.org/10.6092/JOSER_2016_07_01_P55.
17. The Apache Software Foundation. Apache Tinkerpop. 2021. <http://tinkerpop.apache.org/gremlin.html>; Accessed: 22 Feb 2021.
18. Gregor K, Rezende DJ, Wierstra D. Variational intrinsic control. *ArXiv.* 2016. abs/1611.07507. [1611.07507](https://arxiv.org/abs/1611.07507).
19. Eysenbach B, Gupta A, Ibarz J, Levine S. Diversity is all you need: Learning skills without a reward function. *ArXiv.* 2018. abs/1802.06070. [1802.06070](https://arxiv.org/abs/1802.06070).
20. Achiam J, Edwards HA, Amodei D, Abbeel P. Variational option discovery algorithms. *ArXiv.* 2018. abs/1807.10299.
21. Pathak D, Gandhi D, Gupta A. Self-supervised exploration via disagreement. *Proceedings of the 36th International Conference on Machine Learning, in Proceedings of Machine Learning Research* 97:5062–5071. Available from <http://proceedings.mlr.press/v97/pathak19a.html>.
22. Aubret A, Matignon L, Hassas S. A survey on intrinsic motivation in reinforcement learning. *ArXiv.* 2019. abs/1908.06976. [1908.06976](https://arxiv.org/abs/1908.06976).
23. Lehman J, Stanley KO. Abandoning objectives: Evolution through the search for novelty alone. *Evol Comput.* 2011;19(2):189–223. https://doi.org/10.1162/EVCO_a_00025.
24. Cully A, Demiris Y. Quality and diversity optimization: A unifying modular framework. *IEEE Trans Evol Comput.* 2018;22(2):245–59. <https://doi.org/10.1109/TEVC.2017.2704781>.
25. Kim S, Coninx A, Doncieux S. From exploration to control: learning object manipulation skills through novelty search and local adaptation. *arXiv.* 2019. abs/1901.00811. [1901.00811](https://arxiv.org/abs/1901.00811).
26. Cully A. Autonomous skill discovery with quality-diversity and unsupervised descriptors. In: *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '19.* New York: ACM; 2019. p. 81–89. <https://doi.org/10.1145/3321707.3321804>.
27. Kumar S, Renaudin V, Aoustin Y, Le-Carpentier E, Combettes C. Model-based and experimental analysis of the symmetry in human walking in different device carrying modes. In: 2016 6th IEEE International Conference on Biomedical Robotics and Biomechanics (BioRob); 2016. p. 1172–9. <https://doi.org/10.1109/biorob.2016.7523790>.
28. Schaal S. In: Kimura H, Tsuchiya K, Ishiguro A, Witte H, editors. *Dynamic Movement Primitives - A Framework for Motor Control in Humans and Humanoid Robotics.* Tokyo: Springer; 2006, pp. 261–80.
29. Langosz M. *Evolutionary Legged Robotics.* Germany: Doctoral dissertation, University of Bremen; 2018.
30. Ha D, Schmidhuber J. World models. *arXiv.* 2018. abs/1803.10122.
31. Haaroja T, Hartikainen K, Abbeel P, Levine S. Latent space policies for hierarchical reinforcement learning. *arXiv.* 2018. abs/1804.02808.
32. Florensa C, Duan Y, Abbeel P. Stochastic neural networks for hierarchical reinforcement learning. *arXiv.* 2017. abs/1704.03012.
33. Lynch C, Khansari M, Xiao T, Kumar V, Tompson J, Levine S, Sermanet P. Learning latent plans from play. In: *Conference on Robot Learning*; 2020. p. 1113–32.
34. Higgins I, Matthey L, Glorot X, Pal A, Uria B, Blundell C, Mohamed S, Lerchner A. Early visual concept learning with unsupervised deep learning. *arXiv.* 2016. abs/1606.05579.
35. Plaut DC, Hinton GE. Learning sets of filters using back-propagation. *Comput Speech Language.* 1987;2(1):35–61.
36. Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. *Science.* 2006;313(5786):504–7.
37. Chen N, Karl M, Van Der Smagt P. Dynamic movement primitives in latent space of time-dependent variational autoencoders. In: 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). IEEE; 2016. p. 629–36. <https://doi.org/10.1109/humanoids.2016.7803340>.
38. Wang Z, Merel J, Reed S, Wayne G, de Freitas N, Heess N. Robust imitation of diverse behaviors. 2017. *arXiv preprint arXiv:1707.02747*.
39. Dilokthanakul N, Mediano PA, Garnelo M, Lee MC, Salimbeni H, Arulkumaran K, Shanahan M. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv.* 2016. abs/1611.02648.
40. Chiesa M. *Radical Behaviorism: The Philosophy and the Science.* Authors Cooperative, Inc; 1994.
41. Gutzeit L, Kirchner EA. Automatic detection and recognition of human movement patterns in manipulation tasks. In: *PhyCS*; 2016. p. 54–63.
42. Gutzeit L, Fabisch A, Petzoldt C, Wiese H, Kirchner F. Automated Robot Skill Learning from Demonstration for Various Robot Systems. In: Benz Müller C, Stuckenschmidt H, editors. *KI 2019: Advances in Artificial Intelligence.* KI 2019. Lecture Notes in Computer Science, vol 11793. Cham: Springer; 2019. https://doi.org/10.1007/978-3-030-30179-8_14.
43. Leohold S. *Active Reward Learning* Fuer Master Thesis; University Bremen; 2019 Gesteht.
44. Kim SK, Kirchner EA, Stefes A, Kirchner F. Intrinsic interactive reinforcement learning - Using error-related potentials for real world human-robot interaction. *Sci Rep.* 2017;7:17562. <https://doi.org/10.1038/s41598-017-17682-7>.
45. Tenorth M, Beetz M. Knowrob – knowledge processing for autonomous personal robots. In: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE; 2009. p. 4261–6.
46. Beetz M, Mösenlechner L, Tenorth M. CRAM: A Cognitive Robot Abstract Machine for everyday manipulation in human environments. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE; 2010. p. 1012–7. <https://doi.org/10.1109/IROS.2010.5650146>. [http://ieeexplore.ieee.org/document/5650146/](http://ieeexplore.ieee.org/document/5650146).
47. Stock S, Mansouri M, Pecora F, Hertzberg J. Online task merging with a hierarchical hybrid task planner for mobile service robots. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE; 2015. p. 6459–64. <https://doi.org/10.1109/IROS.2015.7354300>.
48. Bercher P, Keen S, Biundo S. Hybrid planning heuristics based on task decomposition graphs. In: *Proceedings of the 7th Annual Symposium on Combinatorial Search, SoCS 2014, vol. 2014-Janua*; 2014. p. 35–43.

49. Nau D, Au T.-c., Ilghami O, Kuter U, Murdock JW, Wu D, Yaman F. SHOP2: An HTN Planning System. *Syst Res.* 2003;20:379–404.
50. Fox M, Long D. PDDL2. 1: An Extension to PDDL for Expressing Temporal Planning Domains. *J Artif Intell Res (JAIR).* 2003. <https://doi.org/10.1613/jair.1129>.
51. Höller D, Behnke G, Bercher P, Biundo S, Fiorino H, Pellier D, Alford R. HDDL : An Extension to PDDL for Expressing Hierarchical Planning Problems. In: *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI 2020)*. AAAI Press; 2020. https://www.uni-ulm.de/fileadmin/website_uni_ulm/iui.inst.090/Publikationen/2020/Hoeller2020HDDL.pdf.
52. Dornhege C, Eyerich P, Keller T, Trüg S, Brenner M, Nebel B. Semantic Attachments for Domain-Independent Planning Systems. In: Prassler E, editor. *Towards Service Robots for Everyday Environments*. Springer Tracts in Advanced Robotics, 2012, vol 76. Berlin: Springer. https://doi.org/10.1007/978-3-642-25116-0_9.
53. Roehr TM. Automated Operation of a Reconfigurable Multi-Robot System for Planetary Space Missions. PhD thesis, University Bremen. 2019.
54. Kumar S, Szadkowski K. A. v., Mueller A, Kirchner F. An Analytical and Modular Software Workbench for Solving Kinematics and Dynamics of Series-Parallel Hybrid Robots. *J Mech Robot.* 2020;12(2):. <https://doi.org/10.1115/1.4045941.021114>. https://asmedigitalcollection.asme.org/mechanismsrobotics/article-pdf/12/2/021114/6481776/jmr_12_2_021114.pdf.
55. von Szadkowski K, Reichel S. Phobos: A tool for creating complex robot models. *J Open Source Softw.* 2020;5(45):1326. <https://doi.org/10.21105/joss.01326>.
56. Lloyd S. Least squares quantization in pcm. *IEEE Trans Inform Theory.* 1982;28(2):129–37.
57. Dinh L, Sohl-Dickstein J, Bengio S. Density estimation using real nvp. 2016. arXiv preprint arXiv:1605.08803.
58. Chen TQ, Rubanova Y, Bettencourt J, Duvenaud DK. Neural ordinary differential equations. In: *Advances in Neural Information Processing Systems*; 2018. p. 6571–83.
59. Langosz M. MARS (Machina Arte Robotum Simulans). GitHub. 2021. <https://github.com/rock-simulation/mars>.
60. Kumar S, Wöhrle H, de Gea Fernández J, Müller A, Kirchner F. A survey on modularity and distributivity in series-parallel hybrid robots. *Mechatronics.* 2020;68:102367. <https://doi.org/10.1016/j.mechatronics.2020.102367>.
61. Li Y, Akbar S, Oliva J. ACFlow: Flow Models for Arbitrary Conditional Likelihoods. *Proceedings of the 37th International Conference on Machine Learning*. In: *Proceedings of Machine Learning Research* 119:5831–5841. Available from <http://proceedings.mlr.press/v119/li20a.html>.
62. Holzinger A, Plass M, Kickmeier-Rust M, Holzinger K, Crişan GC, Pintea C-M, Palade V. Interactive machine learning: experimental evidence for the human in the algorithmic loop. *Appl Intell.* 2019;49(7):2401–14.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
