Using Knowledge Representation and Task Planning for Robot-agnostic Skills on the Example of Contact-Rich Wiping Tasks

Matthias Mayr[1], Faseeh Ahmad[1], Alexander Duerr[1] and Volker Krueger[1]

Abstract— **The transition to agile manufacturing, Industry 4.0, and high-mix-low-volume tasks require robot programming solutions that are flexible. However, most deployed robot solutions are still statically programmed and use stiff position control, which limit their usefulness.**

In this paper, we show how a single robot skill that utilizes knowledge representation, task planning, and automatic selection of skill implementations based on the input parameters can be executed in different contexts. We demonstrate how the skill-based control platform enables this with contact-rich wiping tasks on different robot systems. To achieve that in this case study, our approach needs to address different kinematics, gripper types, vendors, and fundamentally different control interfaces. We conducted the experiments with a mobile platform that has a *Universal Robots UR5e* 6 degree-of-freedom robot arm with position control and a 7 degree-of-freedom *KUKA iiwa* with torque control.

I. INTRODUCTION

The need for flexible and skill-based robot control systems in industrial robot applications is becoming increasingly important. With the rise of automation and the growing complexity of tasks, robots must be able to adapt to changing environments and perform a variety of tasks without lengthy configuration times. This requires a control system that can respond quickly and accurately to changes in the environment and the task at hand. However, most of the deployed systems are statically programmed and perform repetitive tasks. This introduces the need for a flexible and skill-based robot control system in industrial robot applications that provides flexibility and supports agile manufacturing.

First platforms for skill-based systems such as *ClaraTy* [1] or *LAAS* [2], paved the way. In the area of knowledge integration frameworks, the system built around the *Rosetta ontology* [3], [4] and *Knowrob* [5], [6] created a solid foundation. The latter addresses the service domain, while the former is placed in the context of industrial robotics. To transfer skills from one robot to another is nontrivial if robots of different types and vendors are involved. For instance, [7] demonstrated skill transfer for contact-free motions within a single family of robots that share the same vendor-supplied high-level interface *ABB RAPID*. In that paper, the simulated experiments were conducted with high-gain position control. This is, of course, still a standard practice within many tasks that have little variation and high certainty. However, the increasing number of contact-rich tasks that manipulators should be able to solve [8] indicates

[1]Department of Computer Science, Faculty of Engineering (LTH), Lund University, SE 221 00 Lund, Sweden. E-mail: <firstname>.<lastname>@cs.lth.se.
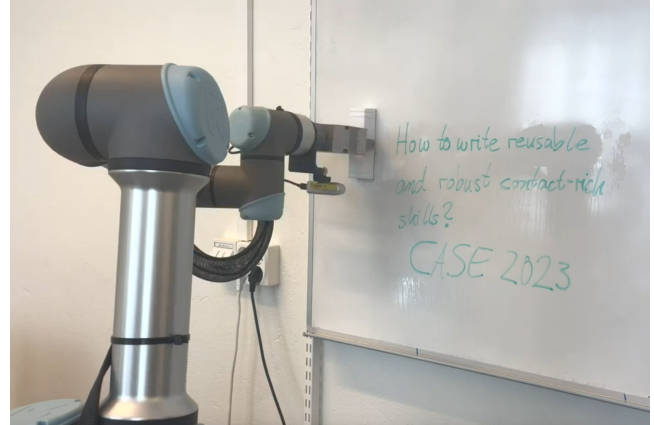
Fig. 1. With the proposed abstractions and knowledge representation, several different robot systems can execute the same skills while still using their respective hardware such as grippers and controllers.

a clear need for flexible and compliant systems. Achieving compliant behaviors heavily depends on the hardware and vendor, and there are no standard interfaces available for this purpose. In this paper, we explore the question of how skills can be transferred between collaborative robots, such as the *Universal Robot* robots and the *KUKA iiwa* or the *Franka Emika Robot (Panda)*. To do this, we develop a suitable representation of the available robot hardware, a knowledge infrastructure to maintain this hardware knowledge, and a suitable compliant controller that works across different types of robot arms. To validate and verify our claims, we test our architecture on a whiteboard wiping task, as shown in Fig. 1, using a *UR5e* and a *KUKA iiwa*. The whiteboard wiping task is an example of a typical contact-rich task that requires a compliant controller. For this, the *UR5e* has a force-torque sensor on the wrist, while robots such as the *iiwa* and the *Panda* have torque sensors in their joints.

In summary, our paper makes the following contributions:

1) We present a complete architecture that represents, plans, and executes skills such that they can be transfered between robots.
2) We demonstrate these ideas based on a whiteboard wiping task using a *UR5e* and a *KUKA iiwa* robot.
3) All essential building blocks, including the impedance controller implementations are available as open-source software.
4) The code is written in a robot-agnostic way, making the results and insights scalable to diverse collaborative robotic hardware and tasks.

II. RELATED WORK

A. Robot Skill Systems

We can classify skill-based architectures based on their target domain. Projects like *CAST* [9] and *KnowRob* [5], [6] are for service robotics and general platforms are *ClaraTy* [1], *LAAS* [2] and *SmartMDSD* [10]. However, platforms that address the specific needs of industrial robotics are the most relevant for this work. *CoStar* [11] has shown to integrate multiple robot systems and provides a graphical user interface (GUI) for the design and the monitoring of skills. However, it lacks explicit knowledge representation and the possibility to perform task planning. The use of ontologies to describe skills in robotic applications has been explored in various areas, including medical and surgical robotics [12], [13], perception ontologies [14] for robots, and more [15]. Among these, [3], [4], [7] present a knowledge-based approach for programming synchronized motions between robotic systems and human-robot interactions that is particularly relevant for our work.

In [7], the authors demonstrate the effectiveness of their approach through experiments that involve transformations of dual arm programs and the transfer of skills between robot systems with different kinematics. Their approach uses finite state machines to implement skill behavior and generates *ABB RAPID* code that is specific to *ABB* robot controllers. Our work differs in that we use behavior trees for skill behavior implementation and our implementation is not limited to a particular robot controller. In addition, [7] was validated only in simulation, and the tasks were not contact-rich.

B. Robot Motions

In the literature, Movement Primitives (MPs) have been a reliable tool for generating robot motions. Dynamic Movement Primitives (DMPs) [16], [17] and Probabilistic Movement Primitives (ProMPs) [18] have been successfully used to generate arm motions. Both DMPs and a motion representation based on behavior trees and motion generators (BTMG) that will be introduced in Section IV-A use attractor landscapes to reach a target location. DMPs are non-linear dynamical systems that allow trajectory control, and, like our BTMGs, they have a non-linear forcing term that enables the modification of the original trajectory. However, in contrast to DMPs, the parameters of BTMGs can be specified explicitly, giving us a better understanding of the robot motion and allowing for a better intregration with knowledge representation.

Wiping Applications: In [19], [20], [21] the authors discuss research on robotic manipulation of wiping tasks using artificial intelligence reasoning methods. [19] proposes a classification of compliant manipulation tasks based on symbolic effects, subcategorizes wiping tasks, and demonstrates how to concretize actions based on the example of shards sweeping with a broom. [20] investigates the reasoning and action execution problems involved in the execution of wiping tasks and proposes a high-level abstraction
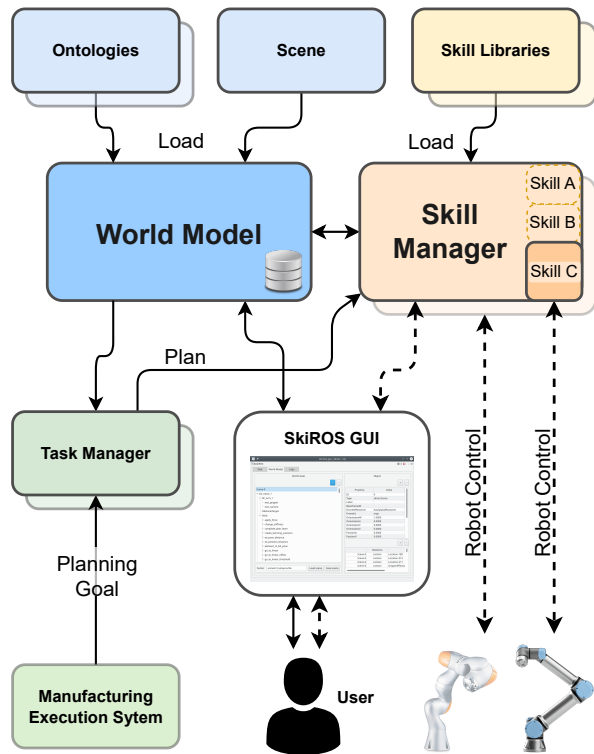

Fig. 2. An outline of the SkiROS2 architecture. The world model stores the knowledge about the relations, environment and the skills. The skill manager loads and executes the skills. Dashed lines show control flows and solid lines show information flows. Shaded blocks indicate possible multiple instances.

representation of wiping tasks to develop generalized action execution mechanisms. [21] combines reasoning methods and compliant robotic manipulation to solve wiping tasks by introducing a qualitative particle distribution model, an approach to generate whole-body wiping motions based on effect-oriented policies, an approach to assess wiping motion quality, and the semantic interpretation of contact situations. Rather than achieving the best wiping skill, the focus of our paper is on presenting a complete architecture that represents, plans, and executes skills, and then using wiping as a challenging example to demonstrate these ideas.

III. TRANSFERABLE ROBOT SKILLS

In this section we introduce the prerequisites that enable to transfer skills between different robot systems. In addition to the skill platform *SkiROS2* and its skill model, the necessary aspects of knowledge presentation and task planning are introduced. The control implementations and motion generation are introduced in Section IV.

A. SkiROS2

We utilize the skill-based robot control platform *SkiROS2* [22] to implement transferable robot skills. It is used in several research projects for motion generation [23], robot coordination [8] and task-level planning with reinforcement learning (RL) [24], [25], [26], [27]. The architecture is shown in Fig. 2 and the main components are the skill

manager and the world model. The world model (WM) contains the knowledge about the world and represents it in a Resource Description Framework (RDF) database. Ontologies such as the *Suggested Upper Merged Ontology (SUMO)* describe the available concepts such as objects or skills and their relations to each other. In addition to that, a *scene* contains the concrete instances of the concepts that are defined in the ontologies. Such instances can be the present robot systems, workstations or objects that are currently known.

The skill manager is responsible for loading a set of specified skills from the skill libraries. Each skill that is successfully loaded gets a semantic description in the WM. This includes the parameters and pre-, hold- and post-conditions. In addition to that, the skill manager provides services to start, pause and stop skills and to manage their execution.

The task manager accepts the planning goals in Planning Domain Definition Language (PDDL) [28], [22]. Such goals can come from a user or an external entity such as a manufacturing execution system (MES). An example of such a goal is to place an object at a specific location: `(skiros:contain skiros:Location-1 skiros:Product-1)`. The task manager automatically creates a PDDL planning domain from the knowledge in the WM. The currently available skills are added with their preconditions and effects, and the relevant instances in the scene are automatically added to a planning problem. Then a PDDL planner is called to find a sequence of skills that is guaranteed to be optimal [29]. Finally, the plan is converted back into a behavior tree (BT), automatically expanded with the relevant skill implementations and grounded with the instances in the WM [22].

B. Skill Model

We define a skill as a parametric procedure that changes the world from some initial state to some new state [30]. It is shown in Fig. 3. A skill consists of two main components: a skill description and the skill implementation. A skill description contains the skill parameters and the necessary pre-, hold- and post-conditions. These define what prerequisites need to exist to execute a skill and state the expected effects of running it. In the task manager, these conditions can be used for automatic task planning. Furthermore, the parameters state the input and output of a skill. There are three different types of parameters: 1) required, 2) optional and 3) inferred. An example of a required parameter of a "pick" skill is to specify the robot arm, while an optional parameter could be to slow down the execution for testing. The inferred parameters are reasoned about when starting the skill. An example is a "gripper" parameter for a pick skill that already has the arm as a required parameter. With a precondition we can define a relation between both of them and we can use the knowledge in the WM to specify the inferred parameter automatically [22].

In addition to the skill description that describes the semantic actions of a skill, the skill implementation is a
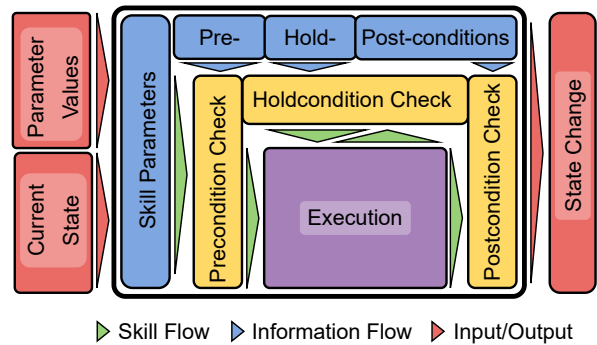

Fig. 3. The conceptual model of a skill in SkiROS2. Pre-, and hold-conditions ensure that the skill is only executed in the correct world state. Post-conditions check if the desired changes have been achieved.

concrete version for the execution. Each skill implementation implements exactly one description with its parameters and pre-, hold- and post-conditions. There are two complexities of skills: Primitive skills that are semantically atomic and compound skills. Primitive skills provide a Python interface for the initialization, start, execution and stop of a skill. In contrast to that, compound skills allow to connect an arbitrary number of primitive skills and compound skills in a BT [31], [22]. A skill implementation is also allowed to change the skill description, such as adding addition conditions or modifying existing conditions. This is important to support multiple, specific skill implementations. Examples of this are gripper actuation skills. A basic skill description can have a parameter "gripper" of the WM type `Element("rparts:GripperEffector")` and a Boolean parameter "opening state". The implementation of such a skill depends on the actual gripper model and a specific implementation can set the parameter to a specific subtype of `rparts:GripperEffector`, such as `scalable:RobotiqGripper`.

C. Knowledge Representation and Task Planning

An important aspect is the separation of skills and tasks. Skills are written in a parametric way to be used in various applications and different tasks. This separation is achieved by explicitly storing knowledge in the WM. Since the conditions and task-specific knowledge is not implicitly represented in the skills, this allows for a modular and extensible design that easily allows to add scenarios or robot systems at a later point. The WM stores this knowledge in a semantic RDF database that utilizes the Web Ontology Language (OWL). SkiROS2 introduces its own ontology, which contains the necessary concepts and relations for skill modeling and reasoning. It is based on the *Core Ontology for Robotics and Automation (CORA)* that is in the *IEEE Standard Ontologies for Robotics and Automation (IEEE Std 1872™-2015)*.

An example for a part of the semantic description of a robot arm is in Listing 1. The listing contains the relevant knowledge for the trajectory generation and motion execution. This includes, but is not limited to properties

```
scalable:Ur5−2 a scalable:Ur5, owl:NamedIndividual ;
    rdfs:label "scalable:ur5"^^xsd:string ;
    skiros:BaseFrameId "cora:Robot−1"^^xsd:string ;
    skiros:CartesianGoalAction "/cartesian_trajectory_generator/goal_action"^^xsd:
        string ;
    skiros:OverlayMotionService "/cartesian_trajectory_generator/overlay_motion"
        ^^xsd:string ;
    skiros:CartesianStiffnessTopic "/cartesian_param_filter/stiffness_goal"^^xsd:
        string ;
    skiros:CartesianWrenchTopic "/cartesian_param_filter/force_goal"^^xsd:string ;
    skiros:CompliantController "cartesian_compliance_controller"^^xsd:string ;
    skiros:JointConfigurationController "scaled_pos_traj_controller"^^xsd:string ;
    skiros:DiscreteReasoner "AauSpatialReasoner"^^xsd:string ;
    skiros:FrameId "scalable:Ur5−2"^^xsd:string ;
    skiros:LinkedToFrameId "ur5e_base_link"^^xsd:string ;
    skiros:MotionExe "/scaled_pos_traj_controller/follow_joint_trajectory"^^xsd:
        string ;
    skiros:MoveItGroup "manipulator"^^xsd:string ;
    skiros:MoveItReferenceFrame "ur5e_base_link"^^xsd:string ;
    skiros:MoveItTCPLink "ur5e_tcp_link"^^xsd:string ;
    skiros:hasA scalable:WsgGripper−3 .
```

such as Robot Operating System (ROS) interfaces of a robot and the relevant aspects of a motion planning configuration. In addition there are elements needed for reasoning about relations such as the attached gripper.

Tasks are executed by the skill manager and there are two ways to start them: 1) Select and parameterize a skill manually. For example in the GUI or from an application programming interface (API) or 2) by providing a planning goal to the task manager and executing the skill sequence. Tasks are concrete skill (sequence) instances with a specific goal. Within a task, the skills can utilize a blackboard to exchange information with their parameters.

IV. CONTROL AND MOTION GENERATION

A. Behavior Trees and Motion Generators

Behavior Trees (BT) are mathematical models used for plan representation and execution. They have been shown to perform reliably in games [32], artificial intelligence and robotics [33]. A BT is an acyclic graph defined with nodes and edges defined in a typical parent-child relationship. The nodes are of two types: 1) *control flow* and 2) *execution*. The *control flow* nodes are responsible for directing the flow of the tree and are traditionally classified on the basis of their operation as: 1) sequence, 2) selector, 3) parallel and 4) decorator [31], [34]. In essence, a BT works by means of sending a tick signal from the *Root* node. The signal then moves through the nodes controlled by the *control flow* nodes. The return statements from the *execution* nodes are *running*, *success* or *failure*. In order to combine BTs with task-level planning, [22] proposes *extended BTs*.

A motion generator (MG) [23] generates an arm motion using impedance controllers to control the end effector (EE) in Cartesian space. A MG follows a simple attractor landscape that uses a virtual spring to attract the EE to the desired location, whereas virtual dampers allow safe motion by slowing down the overall motion. In addition, it also allows for a deviation by superimposing a motion on top of the original trajectory. For example, in [25] we used a Cartesian linear trajectory superimposed by an Archimedes spiral to solve the peg-in-hole task.

A Behavior Tree and Movement Generator (BTMG) combines the strengths of both BT and MG to model robotic skills. A BTMG is a parameterized representation that allows us to specify not only the structure of BT but also the properties of the MG. In essence, the BTMG representation not only covers the aspects of plan representation and execution, but also determines the arm motions for the execution strategies [35].

B. Trajectory Generation

The concept of the MG has been introduced in [23], which consists of both a trajectory generation method and a corresponding controller to execute the generated trajectory on the manipulator. In the following section, we provide a description of the trajectory generation process adopted for this project.

In many industrial settings, operations are performed in a Cartesian task space rather than in the joint space of the robot. Therefore, we argue that motion references are best generated and provided in the Cartesian end-effector space. In this work, we adopt Cartesian linear paths[1] to generate the corresponding trajectories. This is achieved by selecting appropriate acceleration profiles and setting the maximum Cartesian translational and rotational velocities. Furthermore, we offer the option to synchronize motions that have both translational and rotational components. Specifically, the trajectory generator takes the new goal pose as input and outputs end-effector pose references, which are then sent to the corresponding controller.

Aside from specifying a new goal pose, it is also possible to overlay additional motions on top of the reference pose. This capability has been utilized in various studies such as [36], [24], [25], where an Archimedes spiral is used to find a hole in an insertion procedure. In this work, we apply a sine motion to improve the wiping performance of the robot.

C. Compliant Control

Manipulation tasks that have uncertainties need compliant control solutions to provide the necessary flexibility. Such uncertainties can come from either the task itself or inaccurate knowledge about the placement of objects. Another source could be the placement of the robot, which becomes relevant when working with a mobile robot. The robot systems used in the case study provide different interfaces for control commands and therefore also have different solutions for compliant control.

Cartesian Impedance Control: Robot systems such as the *KUKA iiwa* or the *Franka Emika Robot* offer an interface to send commanded torques for each of the joints. This can be used to perform Cartesian impedance control [37] which allows compliant control of the end effector in task space.

---

[1]https://github.com/matthias-mayr/cartesian_trajectory_generator

Torque-controlled robots are typically gravity compensated. The rigid-body dynamics of such a system in the joint space is described as $q \in \mathbb{R}^n$:

$$M(q)\ddot{q} + C(q,\dot{q})\dot{q} = \tau_{\mathrm{c}} + \tau^{\mathrm{ext}} \tag{1}$$

where $M(q) \in \mathbb{R}^{n \times n}$ is the generalized inertia matrix, $C(q,\dot{q}) \in \mathbb{R}^{n \times n}$ captures the effects of Coriolis and centripetal forces, $\tau_{\mathrm{c}} \in \mathbb{R}^n$ represents the input torques, and $\tau^{\mathrm{ext}} \in \mathbb{R}^n$ represents the external torques, with n being the number of joints of the robot.

Moreover, the torque signal commanded by this controller to the robot, τ_{c} in Equation (1), is composed of the superposition of three joint-torque signals:

$$\tau_{\mathrm{c}} = \tau_{\mathrm{c}}^{\mathrm{ca}} + \tau_{\mathrm{c}}^{\mathrm{ns}} + \tau_{\mathrm{c}}^{\mathrm{ext}} \tag{2}$$

where

$\tau_{\mathrm{c}}^{\mathrm{ca}}$ is the torque commanded to achieve a Cartesian impedance behavior with respect to a Cartesian pose reference in the m-dimensional task space, $\xi^{\mathrm{D}} \in \mathbb{R}^m$, in the frame of the end effector of the robot:

$$\tau_{\mathrm{c}}^{\mathrm{ca}} = J^{\mathrm{T}}(q)\left[-K^{\mathrm{ca}}\Delta\xi - D^{\mathrm{ca}}J(q)\dot{q}\right] \tag{3}$$

with $J(q) \in \mathbb{R}^{m \times n}$ being the Jacobian relative to the end-effector (task) frame of the robot, and $K^{\mathrm{ca}} \in \mathbb{R}^{m \times m}$ and $D^{\mathrm{ca}} \in \mathbb{R}^{m \times m}$ being the virtual Cartesian stiffness and damping matrices, respectively.

$\tau_{\mathrm{c}}^{\mathrm{ns}}$ is the torque commanded to achieve a joint impedance behavior with respect to a desired configuration and projected in the null-space of the robot's Jacobian, to not affect the Cartesian motion of the robot's end effector.

$\tau_{\mathrm{c}}^{\mathrm{ext}}$ is the torque commanded to achieve the desired external force command in the frame of the end effector of the robot, $F_{\mathrm{c}}^{\mathrm{ext}}$:

$$\tau_{\mathrm{c}}^{\mathrm{ext}} = J^{\mathrm{T}}(q)F_{\mathrm{c}}^{\mathrm{ext}} \tag{4}$$

In this work we utilize $\tau_{\mathrm{c}}^{\mathrm{ext}}$ and $\tau_{\mathrm{c}}^{\mathrm{ca}}$ in combination with changing stiffnesses by passing configurations through skill primitives.

Cartesian Compliance Control: In contrast to the *KUKA iiwa*, the *Universal Robots UR5e* like many other industrial manipulators does not support direct torque control. This means that a compliant control solution is needed that can work with joint position or joint velocity commands. Recently, such a method was introduced through forward dynamics compliance control [38]. It combines admittance, impedance and force control into one control strategy. The control loop is closed only through a force-torque sensor, which makes it system independent. Internally it utilizes a forward dynamics simulation of a virtual model of the robot system to map Cartesian inputs to joint commands. For further details, we refer the interested reader to [38] and [39]. Like the impedance controller that was introduced in the previous section, this controller solution also allows for the runtime modification of applied external forces, reference poses and stiffness changes. We utilize these capabilities when modeling behaviors with BTs.


Fig. 4. The *iiwa* wiping task on a table. The different hardware needs a different controller setup and gripper skills.

V. CASE STUDY: WIPING TASK

We demonstrate the usage of skills in different tasks and on different robot system with the example of a wiping task.

A. Challenges

The task at hand and the robot systems selected for this study present a range of intriguing challenges. Firstly, the task instances involve surfaces with different properties, namely a smooth whiteboard and a rough industrial table. The whiteboard is mounted to a wall, whereas the table surface is parallel to the floor. Moreover, we assume that the precise distance between the robot and the surface may not be known in advance. The robot systems are equipped with different grippers, each of which has its own communication interface. Finally, the robots have distinct kinematics, vendor-supplied programming solutions, and control interfaces. In this context, we demonstrate how a skill that utilizes knowledge representation, planning and the automatic skill implementation selection can be executed in different context on different robot systems.

B. Implementation

The implementation of such a skill incorporates several key aspects that are discussed in [23]. Specifically, we utilize a motion generator (MG) that comprises of both trajectory generation (Sec.IV-B) and a controller (Sec.IV-C).

As inputs for the skill, only the robot arm and a surface to clean need to be selected among the available instances in the WM. All other necessary data, such as robot-specific properties (e.g., the controller to use) and the surface properties (e.g. the dimensions, force to apply) are automatically fetched from the WM as part of the knowledge integration. Apart from the parameters, this skill also contains the necessary pre- and postconditions for planning. The integrated planner in SkiROS can receive a goal such as (`skiros:clean scalable:Workstation-1186 scalable:Cell-12`). For instance, one precondition of the skill is that the tool must already be held in the gripper.
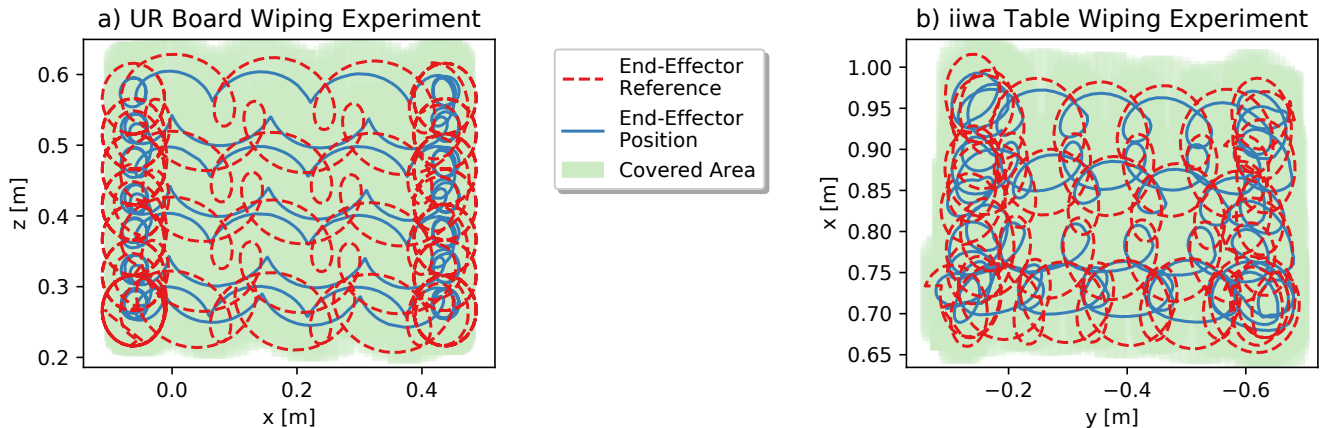
Fig. 5. The resulting reference positions and actual positions of the end effectors in a) the whiteboard wiping task and with the *UR5e* b) the table wiping task with the *iiwa*. The shaded green shows the surface area that has been covered by the eraser that is used for cleaning.

This is leveraged in task-level planning, where the robot is required to pick up the tool first. In case of a mobile system, this may entail driving to a different location. The postcondition of the task is that the surface is cleaned.

To begin the wiping process, the end effector is moved to the corner of the surface to be cleaned. Once there, it disables the stiffness of the compliant controller along the normal of the surface and applies a predetermined force. In our experiments, $8\,\mathrm{N}$ are specified for surfaces that are easier to clean, such as the whiteboard, and $18\,\mathrm{N}$ for the table. The robot then applies an overlay motion, in this case a circular motion, while performing the wiping action. From there the skill executes laps by moving the robot's end-effector to the right, then down, and finally left until the entire surface is covered. Finally, upon completion of the task, the robot stops the force application, increases the stiffness along the surface normal, and moves the end-effector away from the surface.

C. Experiments and Discussion

We conducted the experiments using two different robot systems. The first system is a mobile platform that has a *UR5e* 6 degree-of-freedom robot arm with a *Schunk* two finger gripper. The second system is a 7 degree-of-freedom *KUKA iiwa* with a *Robotiq* 3-finger gripper. We selected two different surfaces for the cleaning task, each with different orientations and surface properties.

The resulting reference path and the actual path for the *UR5e* are shown in Fig. 5a). The implementation spends more time in the left and right extremes since the sub-trajectory ends there and it switches to new lane. This can be adressed by blending between trajectories. During the execution we also noticed subtle vibrations, especially when the arm was stretched out.

In our second experiment we used the *KUKA iiwa* robot system to wipe the surface of a table. As in the previous experiment, the important attributes were represented in the WM, and the skill could be started by simply selecting the surface and the robot arm. The reference path and the actual path for the robot end effector during the task is shown in Figure 5b). The controller is very stable and there are no vibrations. We observed that the eraser was occasionally lifted slightly during the wiping process. This could potentially be addressed by increasing the rotational stiffness or applying more pressure to the surface.

In both tasks, the respective robot systems were able to completely clean the surface with the proposed implementation. In principle, the surface can also be cleaned using only the Cartesian linear motion. However, as intuitively expected, the addition of an additional overlay motion improves the cleaning performance in practice.

VI. CONCLUSIONS AND FUTURE WORK

The transition towards Industry 4.0 brings challenges such as small batch sizes, constantly changing tasks and environments. Together with the transition towards human-robot colaboration and handling contact-rich tasks, it creates the need for platforms that can address these challenges.

In this paper we presented a complete architecture that allows to transfer skills between different robot systems. We explain the necessary prerequisites and knowledge representation. We demonstrated this by successfully performing contact-rich wiping tasks with a *KUKA iiwa* and a *Universal Robots UR5e*. All essential building blocks including the skill-based system and the controllers are available as open-source software and the results are expected to scale to other hardware and tasks as well.

While the wiping performance was sufficient in our tasks, it can potentially be increased by learning a good combination of path velocities and applied force. A visual inspection with a camera is a next inuitive step. We are also looking into an evaluation with other platforms such as the *Franka Emika Robot (Panda)*.

REFERENCES

[1] R. Volpe, I. Nesnas, T. Estlin, D. Mutz, R. Petras, and H. Das, "The claraty architecture for robotic autonomy," in *2001 IEEE Aerospace Conference Proceedings (Cat. No. 01TH8542)*, vol. 1. IEEE, 2001, pp. 1–121.

[2] S. Bensalem, M. Gallien, F. Ingrand, I. Kahloul, and N. Thanh-Hung, "Designing autonomous robots," *IEEE Robotics & Automation Magazine*, vol. 16, no. 1, pp. 67–77, 2009.

[3] M. Stenmark and J. Malec, "Knowledge-Based Industrial Robotics," *Frontiers in Artificial Intelligence and Applications*, pp. 265–274, 2013.

[4] ——, "Knowledge-based instruction of manipulation tasks for industrial robotics," *Robotics and Computer-Integrated Manufacturing*, vol. 33, pp. 56–67, Jun. 2015.

[5] M. Tenorth and M. Beetz, "KNOWROB — knowledge processing for autonomous personal robots," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2009, pp. 4261–4266.

[6] ——, "KnowRob: A knowledge processing infrastructure for cognition-enabled robots," *The International Journal of Robotics Research*, vol. 32, no. 5, pp. 566–590, Apr. 2013.

[7] E. A. Topp, M. Stenmark, A. Ganslandt, A. Svensson, M. Haage, and J. Malec, "Ontology-based knowledge representation for increased skill reusability in industrial robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 5672–5678.

[8] D. Wuthier, F. Rovida, M. Fumagalli, and V. Krüger, "Productive multitasking for industrial robots," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 12 654–12 661.

[9] J. L. Wyatt, A. Aydemir, M. Brenner, M. Hanheide, N. Hawes, P. Jensfelt, M. Kristan, G.-J. M. Kruijff, P. Lison, A. Pronobis *et al.*, "Self-understanding and self-extension: A systems and representational approach," *IEEE Transactions on Autonomous Mental Development*, vol. 2, no. 4, pp. 282–303, 2010.

[10] D. Stampfer, A. Lotz, M. Lutz, and C. Schlegel, "The smartmdsd toolchain: An integrated mdsd workflow and integrated development environment (ide) for robotics software," *Journal of Software Engineering for Robotics (JOSER)*, vol. 7, no. 1, pp. 3–19, 2016.

[11] C. Paxton, A. Hundt, F. Jonathan, K. Guerin, and G. D. Hager, "Costar: Instructing collaborative robots with behavior trees and vision," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 564–571.

[12] B. Bruno, N. Y. Chong, H. Kamide, S. Kanoria, J. Lee, Y. Lim, A. K. Pandey, C. Papadopoulos, I. Papadopoulos, F. Pecora *et al.*, "The caresses eu-japan project: making assistive robots culturally competent," in *Ambient Assisted Living: Italian Forum 2017 8*. Springer, 2019, pp. 151–169.

[13] B. Gibaud, G. Forestier, C. Feldmann, G. Ferrigno, P. Gonçalves, T. Haidegger, C. Julliard, D. Katić, H. Kenngott, L. Maier-Hein *et al.*, "Toward a standard ontology of surgical process models," *International journal of computer assisted radiology and surgery*, vol. 13, pp. 1397–1408, 2018.

[14] H. Azevedo, J. P. R. Belo, and R. A. Romero, "Ontpercept: a perception ontology for robotic systems," in *2018 Latin American Robotic Symposium, 2018 Brazilian Symposium on Robotics (SBR) and 2018 Workshop on Robotics in Education (WRE)*. IEEE, 2018, pp. 469–475.

[15] S. Manzoor, Y. G. Rocha, S.-H. Joo, S.-H. Bae, E.-J. Kim, K.-J. Joo, and T.-Y. Kuc, "Ontology-based knowledge representation in robotic systems: a survey oriented toward applications," *Applied Sciences*, vol. 11, no. 10, p. 4324, 2021.

[16] A. Ijspeert, J. Nakanishi, and S. Schaal, "Trajectory formation for imitation with nonlinear dynamical systems," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180)*, vol. 2, 2001, pp. 752–757 vol.2.

[17] A. J. Ijspeert, J. Nakanishi, and S. Schaal, "Learning rhythmic movements by demonstration using nonlinear oscillators," in *Proceedings of the ieee/rsj int. conference on intelligent robots and systems (iros2002)*, no. CONF, 2002, pp. 958–963.

[18] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing systems*, vol. 26, 2013.

[19] D. Leidner, C. Borst, A. Dietrich, M. Beetz, and A. Albu-Schäffer, "Classifying compliant manipulation tasks for automated planning in robotics," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015, pp. 1769–1776.

[20] D. Leidner, W. Bejjani, A. Albu-Schäffer, and M. Beetz, "Robotic agents representing, reasoning, and executing wiping tasks for daily household chores," *AUTONOMOUS AGENTS AND MULTI-AGENT SYSTEMS*, 2016.

[21] D. Leidner, G. Bartels, W. Bejjani, A. Albu-Schäffer, and M. Beetz, "Cognition-enabled robotic wiping: Representation, planning, execution, and interpretation," *Robotics and Autonomous Systems*, vol. 114, pp. 199–216, 2019.

[22] F. Rovida, B. Grossmann, and V. Krüger, "Extended behavior trees for quick definition of flexible robotic tasks," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6793–6800.

[23] F. Rovida, D. Wuthier, B. Grossmann, M. Fumagalli, and V. Krüger, "Motion generators combined with behavior trees: A novel approach to skill modelling," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 5964–5971.

[24] M. Mayr, F. Ahmad, K. Chatzilygeroudis, L. Nardi, and V. Krueger, "Combining planning, reasoning and reinforcement learning to solve industrial robot tasks," *IROS 2022 Workshop on Trends and Advances in Machine Learning and Automated Reasoning for Intelligent Robots and Systems*, 2022.

[25] ——, "Skill-based multi-objective reinforcement learning of industrial robot tasks with planning and knowledge integration," in *2022 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2022, pp. 1995–2002.

[26] M. Mayr, C. Hvarfner, K. Chatzilygeroudis, L. Nardi, and V. Krueger, "Learning skill-based industrial robot tasks with user priors," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2022, pp. 1485–1492.

[27] F. Ahmad, M. Mayr, and V. Krueger, "Learning to adapt the parameters of behavior trees and motion generators to task variations," 2023.

[28] C. Aeronautiques, A. Howe, C. Knoblock, I. D. McDermott, A. Ram, M. Veloso, D. Weld, D. W. SRI, A. Barrett, D. Christianson *et al.*, "Pddl| the planning domain definition language," *Technical Report, Tech. Rep.*, 1998.

[29] P. Eyerich, R. Mattmüller, and G. Röger, "Using the context-enhanced additive heuristic for temporal and numeric planning," in *Nineteenth International Conference on Automated Planning and Scheduling*, 2009.

[30] S. Bøgh, O. S. Nielsen, M. R. Pedersen, V. Krüger, and O. Madsen, "Does your robot have skills?" in *Proceedings of the 43rd international symposium on robotics*. VDE Verlag GMBH, 2012.

[31] M. Colledanchise and P. Ögren, *Behavior Trees in Robotics and AI: An Introduction*. Chapman & Hall/CRC Press, 2017.

[32] ——, "How Behavior Trees modularize robustness and safety in hybrid systems," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 1482–1488.

[33] M. Iovino, E. Scukins, J. Styrud, P. Ögren, and C. Smith, "A survey of behavior trees in robotics and ai," 2020.

[34] A. Marzinotto, M. Colledanchise, C. Smith, and P. Ögren, "Towards a unified behavior trees framework for robot control," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 5420–5427.

[35] F. Ahmad, M. Mayr, E. A. Topp, J. Malec, and V. Krueger, "Generalizing behavior trees and motion-generator (btmg) policy representation for robotic tasks over scenario parameters," in *2022 IJCAI Planning and Reinforcement Learning Workshop*, 2022.

[36] M. Mayr, K. Chatzilygeroudis, F. Ahmad, L. Nardi, and V. Krueger, "Learning of Parameters in Behavior Trees for Movement Skills," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.

[37] M. Mayr and J. M. Salt-Ducaju, "A c++ implementation of a cartesian impedance controller for robotic manipulators," *arXiv preprint arXiv:2212.11215*, 2022.

[38] S. Scherzinger, A. Roennau, and R. Dillmann, "Forward dynamics compliance control (fdcc): A new approach to cartesian compliance for robotic manipulators," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 4568–4575.

[39] ——, "Inverse kinematics with forward dynamics solvers for sampled motion tracking," in *2019 19th International Conference on Advanced Robotics (ICAR)*. IEEE, 2019, pp. 681–687.