

# Learning and Composing Primitive Skills for Dual-arm Manipulation

Èric Pairet<sup>1,2</sup>[0000–0002–3363–0426], Paola Ardón<sup>1,2</sup>[0000–0002–3026–0706],  
Michael Mistry<sup>1</sup>[0000–0003–3979–922X], and Yvan Petillot<sup>2</sup>[0000–0002–1596–289X]

<sup>1</sup> Institute of Perception, Action and Behaviour, University of Edinburgh, UK  
[eric.pairet@ed.ac.uk](mailto:eric.pairet@ed.ac.uk) [paola.ardon@ed.ac.uk](mailto:paola.ardon@ed.ac.uk) [mmistry@inf.ed.ac.uk](mailto:mmistry@inf.ed.ac.uk)

<sup>2</sup> Engineering and Physical Sciences, Heriot-Watt University, UK  
[y.r.petillot@hw.ac.uk](mailto:y.r.petillot@hw.ac.uk)


**Abstract.** In an attempt to confer robots with complex manipulation capabilities, dual-arm anthropomorphic systems have become an important research topic in the robotics community. Most approaches in the literature rely upon a great understanding of the dynamics underlying the system’s behaviour and yet offer limited autonomous generalisation capabilities. To address these limitations, this work proposes a modelisation for dual-arm manipulators based on dynamic movement primitives laying in two orthogonal spaces. The modularity and learning capabilities of this model are leveraged to formulate a novel end-to-end learning-based framework which (i) learns a library of primitive skills from human demonstrations, and (ii) composes such knowledge simultaneously and sequentially to confront novel scenarios. The feasibility of the proposal is evaluated by teaching the iCub humanoid the basic skills to succeed on simulated dual-arm pick-and-place tasks. The results suggest the learning and generalisation capabilities of the proposed framework extend to autonomously conduct undemonstrated dual-arm manipulation tasks.

**Keywords:** Learning from Demonstration · Humanoid Robots · Model Learning for Control · Dual Arm Manipulation · Autonomous Agents.

## 1 INTRODUCTION

Complex manipulation tasks can be achieved by endowing anthropomorphic robots with dual-arm manipulation capabilities. Bi-manual arrangements extend the systems competences to efficiently perform tasks involving large objects or assembling multi-component elements without external assistance. These systems not only deal with the challenges of single-arm manipulators, such as trajectory planning and environmental interaction, but also require an accurate synchronisation between arms to avoid breaking or exposing the handled object to stress.

Traditional approaches have addressed the aforementioned challenges by means of control and planning-based methods [18]. These methods depend upon an excellent understanding of the exact model underlying the system’s and task’s dynamics, which are commonly approximated to make the calculations computationally tractable [13]. On top of that, some of these methods lack scalability



**Fig. 1.** iCub humanoid learning to contour an obstacle through kinaesthetic guiding (left), and composing multiple skills to conduct a dual-arm pick-and-place task (right).

and generalisation capabilities, involving hand-defining all possible scenarios and actions [5, 3]. All these issues have motivated the use of more natural techniques for robot programming, such as learning by demonstration (LbD), in which a human movement is recorded to be later reproduced by a robot.

Despite the encouraging possibilities offered by adopting human knowledge for robot control, teaching complex systems, such as dual-arm manipulators, to respond and adapt to a broad case of scenarios remains an open challenge. Particularly, it is expected from a dual-arm system to generalise the provided demonstrations to confront novel scenarios in (a) the task space to deal with the changing requirements about trajectory planning and environmental interaction, and (b) the relative space to ensure the essential synchronisation between arms [12]. However, current learning-based architectures in the literature pursuing autonomy and robustness against the dynamic and unpredictable real-world environments are limited to single-arm arrangements [15, 17, 4]. Contrarily, learning-based frameworks for dual-arm robots do not generalise to undemonstrated states, thus being limited to highly controlled scenarios [21, 7, 19].

This paper presents a novel learning-based framework which endows a dual-arm system with a real-time and generalisable method for manipulation in undemonstrated environments (see Figure 1). The framework models a dual-arm manipulator with a set of dynamic movement primitives laying in two orthogonal spaces to tackle the task’s requirements separately from the synchronisation constraints. The modularity of the DMPs is leveraged to (i) create a library of primitive skills from human demonstrations, and (ii) exploit primitive skills simultaneously and sequentially to create complex behaviours. The potential of the proposal is demonstrated in simulation after recording skills with the iCub humanoid through kinaesthetic guiding. The results suggest the proposal’s suitability to endow a dual-arm robot with the necessary learning and generalisation capabilities to autonomously address novel manipulation tasks.

## 2 DUAL-ARM SYSTEM MODELISATION

This paper pursues an end-to-end learning-based framework which endows a dual-arm system with enhanced generalisation capabilities, meets the synchro-

nisation constraints, and is easily programmable by non-robotics-experts. This work addresses all these requirements by means of learnable and composable primitive skills represented as dynamic movement primitives (DMPs) [9]. This section firstly overviews DMPs and its use in the literature. It then introduces the proposed typology of actions in a dual-arm system, which allows leveraging the strengths of a DMP-based modelisation in the dual-arm context.

## 2.1 Dynamic Movement Primitives

DMPs are a versatile tool for modelling and learning complex motions. They describe the dynamics of a primitive skill as a spring-damper system under the effect of a virtual external force called coupling term. This coupling term allows for learning and reproducing any dynamical behaviour, i.e. primitive skill. Importantly, (a) coupling terms can be learnt from human demonstrations, (b) they can be efficiently learned and generated, (c) a unique demonstration is already generalisable, (d) convergence to the goal is guaranteed, and (e) their representation is translation and time-invariant. Because of all these properties, DMPs are adapted to constitute the fundamental building blocks of this work. Next follows an introduction about DMPs and their usage to encode positional and orientational dynamics, and an overview of some coupling terms in the literature.

**Positional Dynamics.** Let the positional state of a one-degree of freedom (DoF) system be defined by its position, linear velocity and acceleration. Then, the system’s state transition is defined with non-linear differential equations as:

$$\tau \dot{z} = \alpha_x (\beta_x (g_x - x) - z) + f_x(\cdot), \quad (1)$$

$$\tau \dot{x} = z, \quad (2)$$

where  $\tau$  is a scaling factor for time,  $x$  is the system’s position,  $z$  and  $\dot{z}$  respectively are the scaled velocity and acceleration,  $\alpha_x$  and  $\beta_x$  are constants defining the positional system’s dynamics,  $g_x$  is the model’s attractor, and  $f_x(\cdot)$  is the coupling term. The coupling term applying at multiple DoFs at once is defined as  $\mathbf{f}_x(\cdot)$ . The system will converge to  $g_x$  with critically damped dynamics and null velocity when  $\tau > 0$ ,  $\alpha_x > 0$ ,  $\beta_x > 0$  and  $\beta_x = \alpha_x/4$  [9].

**Orientalional Dynamics.** A possible representation of orientations is the unit quaternion  $\mathbf{q} \in \mathbb{R}^4 = \mathbb{S}^3$  [20]. They encode orientations of a system as a whole, thus ensuring the stability of the orientational dynamics integration. Let the current orientational state of a system be defined by its orientation, angular velocity and acceleration. Then, the orientational state transition is described by the following non-linear differential equations:

$$\tau \dot{\boldsymbol{\eta}} = \alpha_q (\beta_q 2 \log(\mathbf{g}_q * \bar{\mathbf{q}}) - \boldsymbol{\eta}) + \mathbf{f}_q(\cdot), \quad (3)$$

$$\tau \dot{\mathbf{q}} = \frac{1}{2} \boldsymbol{\eta} * \mathbf{q}, \quad (4)$$

where  $\mathbf{q}$  is the system’s orientation,  $\boldsymbol{\eta}$  and  $\dot{\boldsymbol{\eta}}$  respectively are the scaled angular velocity and acceleration,  $\alpha_q$  and  $\beta_q$  are constants defining the system’s dynamics,  $\mathbf{g}_q \in \mathbb{S}^3$  is the model’s attractor, and  $\mathbf{f}_q(\cdot) \in \mathbb{R}^4$  is the coupling term. The operators  $\log(\cdot)$ ,  $*$ , and  $\bar{\mathbf{q}}$  denote the logarithm, multiplication and conjugate operations for quaternions, respectively.

**Coupling Terms.** Coupling terms describe the system’s behaviour, thus being useful to learn and retrieve any primitive skill. They are commonly used to encode the positional [9] and orientational [20] dynamics of a motion. Coupling terms are modelled in each dimension as a weighted linear combination of non-linear radial basis functions (RBFs) distributed along the trajectory. Thus, learning a certain movement relies on finding the weights of the RBFs which closely reproduce a demonstrated skill.

More complex behaviours may be achieved by exploiting an additional coupling term simultaneously with the motion-encoding one. This approach has been used to avoid joint limits and constraining the robot’s workspace via repulsive forces pushing the system away from these limits [6]. Coupling terms have also been leveraged for obstacle avoidance with an analytic biologically-inspired approach describing how humans steer around obstacles [8, 17]. Another use is for environmental and self-interaction purposes by means of a controller tracking a desired force profile [7]. To the best of the authors’ knowledge, the practice of using coupling terms simultaneously has been limited to two primitive skills acting on the same frame or space [15]. Contrarily, this work further exploits the DMP modularity to describe a dual-arm system in two orthogonal spaces with the purpose of facing complex scenarios by composing multiple coupling terms.

## 2.2 Dual-arm Primitive Skills Taxonomy

Skills for single-arm manipulation have been well analysed in the robotics community. While some of this knowledge can be extrapolated for a dual-arm manipulator as a whole, their complexity resides in the arms interaction. In the context of manipulation via a dual-arm system, a possible classification of any primitive skill falls into two groups: (a) absolute skills, which imply a change of configuration of the manipulated object in the Cartesian or absolute space  $\mathcal{S}_a$ , e.g. move or turn an object in a particular manner, and (b) relative skills, which exert an action on the manipulated object in the object or relative space  $\mathcal{S}_r$ , e.g. opening a bottle’s screw cap, or hold a parcel employing force contact.

Each type of primitive skill uniquely produces movement in its space since they lay in orthogonal spaces such that  $\mathcal{S}_a \perp \mathcal{S}_r$ . It is natural to expect from a dual-arm system to simultaneously carry out, at least, one absolute and one relative skill to accomplish a task. Let us analyse the task of moving a bottle to a particular position while opening its screw cap. Both end-effectors synchronously move to reach the desired configuration (absolute skill). At the same time, the left end-effector is constrained to hold the bottle upright (relative skill), while the right end-effector unscrews the cap (relative skill).


### 2.3 Dual-arm DMP-based Modelisation

Given the variety of primitive skills that a dual-arm system can execute, this work seeks to model the robotic platform in a generalisable yet modular fashion, which accounts for both absolute and relative skills. To this aim, let us consider the closed kinematic chain depicted in Figure 2 operating in a three-dimensional (3D) workspace  $\mathcal{W} = \mathbb{R}^3 \times \text{SO}(3)$ . Each arm  $i$ , where  $i = \{L, R\}$ , interacts with the same object  $\mathcal{O}$ . In this context, the absolute skill explains the movement of the object  $\mathcal{O}$  in the workspace  $\mathcal{W} = \mathcal{S}_a$ , while the relative skill describes the actions of each end-effector  $i$  in  $\mathcal{S}_r$ , i.e. with respect to the object's reference frame  $\{\mathcal{O}\}$ . Note that  $\{\mathcal{O}\}$  is the centre of the closed-chain dual-arm system.

The state of the closed-chain dual-arm system in the workspace can be described by the position/orientation, linear/angular velocities and accelerations of  $\{\mathcal{O}\}$  in  $\mathcal{S}_a$ . As introduced previously, the system's state transition is subjected to its modelled dynamics. Figure 2 illustrates the proposed modelisation of the system's dynamics in  $\mathcal{S}_a$  as a set of DMPs acting between the object's frame  $\{\mathcal{O}\}$  and its goal configuration  $\mathbf{g}_o$ , which accounts for a desired goal position  $\mathbf{g}_{o_x} \in \mathbb{R}^3$  and orientation  $\mathbf{g}_{o_q} \in \mathbb{R}^4$ . Therefore, three positional DMPs as in (1)-(2) and one orientational DMP as in (3)-(4) are required to encode the system's dynamics in the absolute space  $\mathcal{S}_a = \mathbb{R}^3 \times \text{SO}(3)$ .

In the relative space  $\mathcal{S}_r$ , the dynamics of each end-effector are modelled as DMPs referenced to the object's frame  $\{\mathcal{O}\}$ . Since  $\mathcal{S}_r = \mathbb{R}^3 \times \text{SO}(3)$ , each end-effector dynamic's in the relative frame is described by three positional DMPs as in (1)-(2) and one orientational DMP as in (3)-(4).

Any action referenced to the object's frame can be projected to the end-effectors using the grasping geometry  $\mathbf{G}$  of the manipulated object. This allows computing the required end-effector control commands to achieve a particular absolute task. A detailed explanation of this transformation can be found in [12].



**Fig. 2.** DMP-based modelisation of a closed-chain dual-arm system in the absolute and relative spaces. This model is extended to deal with rotational dynamics.

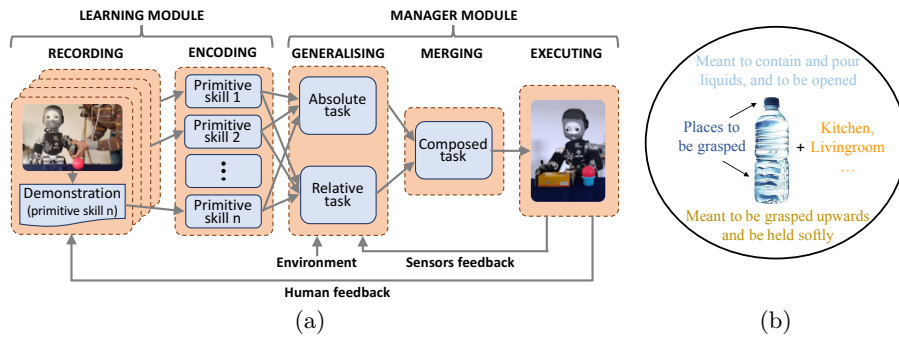
### 3 LEARNING-BASED DUAL-ARM MANIPULATION

To endow a dual-arm manipulator with autonomy and robustness in novel scenarios while being easily programmable and customisable by non-robotics-experts, this work has decomposed and modelled the system’s dynamics and synchronisation constraints as primitive skills lying in the system’s absolute and relative space. Leveraging the formulated modelisation, this work proposes the framework schematised in Figure 3 which creates and manages a library of primitive skills. The framework has two components: (i) a learning module that learns a set of primitive skills from human demonstrations, and (ii) a manager module that combines simultaneously and sequentially these primitives to address a wide range of complex tasks in unfamiliar environments.

#### 3.1 Library Generation

A primitive skill is represented by its coupling term and frame of reference, i.e. either absolute or relative. Learning coupling terms only requires a human demonstrator teaching the characteristic skill. As previously introduced, different coupling terms might be better formulated with different mathematical representations, e.g. a weighted combination of non-linear RBFs to encode the dynamics of a task, an analytical obstacle avoidance expression, or among others, a force profile to control the environmental interaction.

The modularity offered by the proposed DMP-based formulation and its use in two different spaces tackles the hindrance and ambiguity arising when demonstrating all features of a dual-arm task in an all-at-once fashion. This means that instead of learning a task as a whole, the framework harvests a collection of primitive skills. Creating a repertoire of skills referred to as a library, allows the demonstrator to teach in a one-at-a-time fashion, i.e. to focus on one feature of the demonstration at a time [4]. Moreover, this modular library can be



**Fig. 3.** Scheme of the proposed framework. (a) Learning: a library of primitive skills is learnt from human demonstrations. Manager: the primitives are combined simultaneously and sequentially to confront novel environments. (b) The required primitive skills are selected according to the affordance elements of the dual-arm task.

employed for movement recognition purposes, where a demonstrated skill can be compared against the existing ones in the library. If the observed behaviour does not match any existing primitive, it is identified as a new skill and can be added to the framework’s library [10].

### 3.2 Attaching Semantics

The framework needs additional information to successfully conduct a dual-arm manipulation task. Let us consider the robotic task of opening a bottle’s screw cap, where the system needs to select a proper sequence of primitive skills in order to succeed (see Figure 3(b)). This is first a grasping, where each end-effector holds a different component of the bottle, then a synchronous turning referenced in the system’s relative space and finally, a placing and releasing primitives. Therefore, in order to ease this action selection, it is essential to attach a semantic description to each primitive skill.

Semantic labels bridge the gap between the low-level continuous representation of primitives and the high-level description of actions and their influence on objects. An approach to tackle the object affordances challenge consists in combining features from the object and their surroundings to infer on a suitable grasp-action based on their purpose of use [1, 2]. The combination of such elements builds the relationship between context, actions and effects that provide a cognitive reasoning of an object affordance.

### 3.3 Library Management

Each coupling term stored in the framework’s library represents a particular absolute or relative primitive skill. Reproducing a skill consists in using its coupling term as  $\mathbf{f}_x(\cdot)$  or  $\mathbf{f}_q(\cdot)$  in (1)-(4). This computation retrieves the skill’s required accelerations, which can be integrated over time to obtain the skill’s velocities  $\dot{\mathbf{y}}_o$  for an absolute primitive or  $\dot{\mathbf{y}}_{C_i}$  for the end-effector  $i$  relative primitive.

The individual retrieval of primitives already accounts for the inner DMP generalisation capabilities, such as different start and goal configurations, as well as obstacle locations. However, these primitive skills need to be combined to generate more complex movements, such as a pick-and-place task of a bottle accounting for the presence of unexpected obstacles (absolute space), while opening the bottle’s screw cap considering the exerted force (relative space). The presented framework addresses this prerequisite by simultaneously combining different absolute and relative skills as:

$$\begin{bmatrix} \dot{\mathbf{y}}_L \\ \dot{\mathbf{y}}_R \end{bmatrix} = \mathbf{G}^T \sum_{j=1}^J w_j \dot{\mathbf{y}}_{o_j} + \sum_{k=1}^K w_k \begin{bmatrix} \dot{\mathbf{y}}_{C_{L,k}} \\ \dot{\mathbf{y}}_{C_{R,k}} \end{bmatrix}, \quad (5)$$

where  $\dot{\mathbf{y}}_i \in \mathbb{R}^6$  describes the linear and angular velocity commands of the  $i = \{L, R\}$  end-effector satisfying the set of activated primitive skills,  $\mathbf{G} \in \mathbb{R}^{6 \times 12}$  is the global grasp map of the two end-effectors grasp matrices as described in [13],

and  $\dot{\mathbf{y}}_{o_j} \in \mathbb{R}^6$  and  $\dot{\mathbf{y}}_{C_{i,k}} \in \mathbb{R}^6$  are the velocities of the  $j \in [1, J]$  absolute and  $k \in [1, K]$  relative primitive skill stored in the library. Absolute and relative skill selection is conducted with the weights  $w_j$  and  $w_k$ , respectively.

The resulting framework does not only combines skills simultaneously, but also sequentially. This allows the execution of a complex task composed of a sequence of primitives. To do so, a primitive skill is executed by initialising it with the full state (pose, velocities and accelerations) of its predecessor primitive skill. Such an initialisation avoids abrupt jumps in the system’s state.

## 4 RESULTS and EVALUATION

The proposed framework has been evaluated on the iCub humanoid robot. Particularly, the real platform has been used to load the framework’s library with a set of primitive skills learnt from human demonstrations. These skills have been employed in simulation to conduct dual-arm pick-and-place tasks of a parcel in novel scenarios, demonstrating the proposal’s potential for humanoid robots.

### 4.1 Experimental Platform

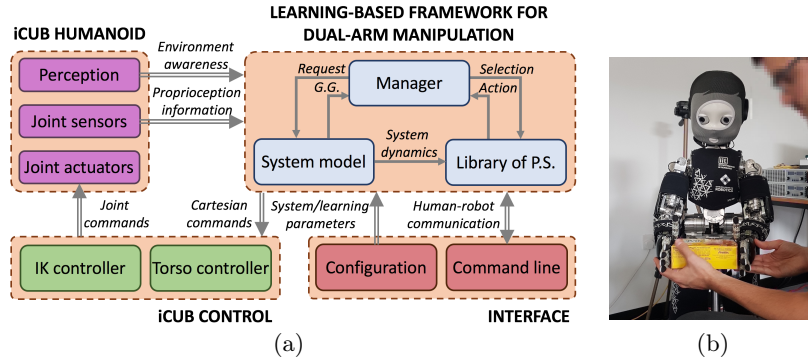
iCub is an open source humanoid robot with 53 DoFs [11] (see Figure 4(b)). The most relevant ones in this work are the three-DoFs on the torso, the two seven-DoFs arms equipped with a torque sensor on the shoulder, and the two nine-DoFs anthropomorphic hands with tactile sensors in the fingertips and palm.

iCub operates under YARP. The deployment of the proposed framework on the iCub platform is schematised in Figure 4. Mainly, four big functional modules can be distinguished: (i) the proposed framework described in this paper (blue blocks), (ii) the real/simulated platform with its visual perception, joint sensors and actuators (magenta blocks), (iii) the end-effectors control via the built-in YARP Cartesian controller [16] and an ad-hoc external torso controller (green blocks), and (iv) the HRI interface to parameterise the desired start and goal configurations for the task, and retrieve the robot’s status (red blocks).

### 4.2 Learning Primitive Skills from Demonstration

For the system to succeed on the dual-arm pick-and-place of a parcel task in novel environments, the framework’s library needs to be loaded with the absolute primitive skills of (i) pick-and-place dynamics on a horizontal surface, (ii) rotational motion around the z-axis, and (iii) obstacle avoidance. Moreover, since the parcel has to be grasped by lateral contact of both end-effectors, the library also requires a relative skill to ensure grasp maintenance, i.e. prevention of contact separation. All these primitive skills have been demonstrated via kineasthetic guiding on the real iCub humanoid robot. To this aim, all joints have been set in gravity compensation, allowing the demonstrator to physically manoeuvre the robot through each primitive. Figure 4(b) depicts the kineasthetic teaching of obstacle avoidance and grasp maintenance primitives.





**Fig. 4.** (a) Layout of the framework deployment on the iCub robot. Note: grasping geometry (GG), primitive skill (PS). (b) iCub humanoid being taught grasp maintenance through kineasthetic guiding.

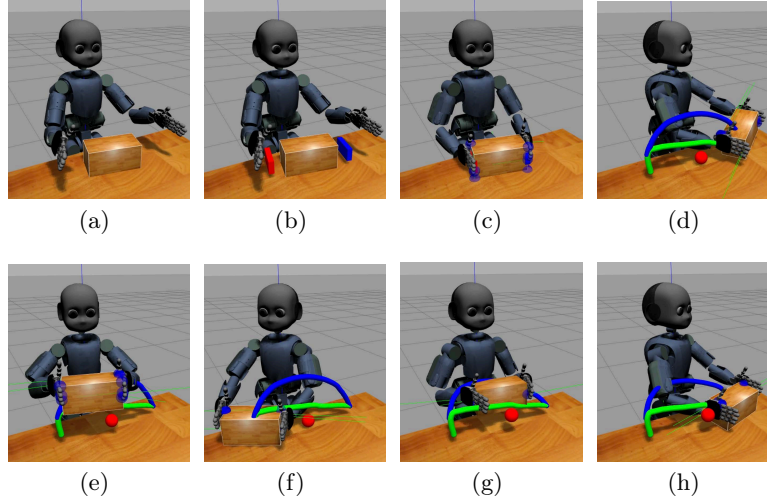
During the demonstrations, proprioception information is retrieved via YARP ports to learn the coupling terms  $\mathbf{f}_x(\cdot)$  and  $\mathbf{f}_q(\cdot)$  in (1)-(4) characterising the different skills. For the pick-and-place and rotational dynamics, the coupling terms are encoded as a weighted linear combination of non-linear RBFs distributed along the trajectory as in [9]. The obstacle avoidance is learnt by finding the best-fitting parameters of the biologically-inspired formulation as in [17]. Finally, the grasp maintenance skill is learnt by setting the parcel's grasping geometry as a pose tracking reference as in [7].

### 4.3 Experiments on Simulated iCub Humanoid

The evaluation of the framework on the pick-and-place setup has been conducted on a simulated iCub robotic platform. Particularly, the four primitive skills previously learnt and loaded in the framework's library are simultaneously and sequentially combined to conduct three consecutive dual-arm pick-and-place task in novel environments (see Figure 5).

Given an initial random configuration laying on the table and within iCub's workspace (see Figure 5(a)), the first action consists of grasping the parcel. This is achieved by retrieving the parcel's configuration, then use the learnt parcel's geometry to compute the grasping points, and finally approach them laterally via the middle-setpoints displayed as red and blue prisms for the right and left end-effector, respectively (see Figure 5(a)-Figure 5(c)). From this stage on, the grasp maintenance skill ensures that both end-effectors are in flat contact with the box to avoid undesired slippage.

The following three consecutive movements require picking-and-placing the parcel between different configurations laying on the central, right and left side of iCub's workspace. The former pick-and-place does not require avoiding any obstacle, thus the built-in DMPs generalisation capabilities are sufficient to address this task (see Figure 5(c)-Figure 5(d)). However, the two latter pick-and-place



**Fig. 5.** iCub humanoid succeeding in novel dual-arm pick-and-place tasks by simultaneously and sequentially combining primitive skills. Demonstrated pick-and-place (green trajectory). Framework’s response (blue trajectory). Obstacle (red sphere). (a) Parcel initial state. (b)-(c) Grasping parcel laterally. (d)-(f) and (f)-(h) Pick-and-place execution with different start, goal and obstacle configurations.

tasks involve adapting the learnt dynamics to address novel scenarios. When the obstacle (red sphere) is collinear with the start and goal positions, i.e. below the demonstrated task (green trajectory), the iCub humanoid circumnavigates the obstacle from the top (see Figure 5(d)-Figure 5(f)). Instead, for an obstacle located forward the demonstration, the framework guides the system through a collision-free trajectory near iCub’s chest (see Figure 5(g)-Figure 5(h)).

The experimental evaluation conducted with the simulated iCub humanoid robot has demonstrated various of the aforementioned framework’s features. Having a repertoire of primitive skills available in the framework’s library allows exploiting them simultaneously and sequentially to confront complex tasks in novel scenarios. The reported case is one of the 16 successful experiments out of a total of 20 trials. In all cases, the robot had to accomplish the three consecutive dual-arm pick-and-place tasks with different start and goal locations, while avoiding novel obstacles and ensuring grasp maintenance. Failure in any of these tasks made the trial unsuccessful. Interestingly, in the four failed trials one of iCub’s forearms collided with the obstacle. This is because the biologically-inspired obstacle avoidance formulation only considers the carried object and should be extended to the object-arm space. The flexibility of the proposed framework could be leveraged to integrate in its library a potential field-inspired approach for obstacle avoidance which also checks for link collisions [14].

## 5 FINAL REMARKS and FUTURE WORK

This work has presented a novel end-to-end learning-based framework which endows a dual-arm manipulator with real-time and generalisable manipulation capabilities. The framework is built upon the proposed extension of the DMP-based modelisation for dual-arm systems, which considers two different frames to reference the movement generation, force interaction and constraints requirements. Based on this arrangement, the proposed framework is twofold: (i) learns from human demonstrations to create a library of primitive skills, and (ii) combines such knowledge simultaneously and sequentially to confront novel scenarios.

The suitability of the proposed approach has been demonstrated in a dual-arm pick-and-place setting, where the iCub humanoid first learnt a repertoire of primitive skills from human demonstrations and then composed such knowledge to successfully generalise to novel scenarios. The framework is not restricted to the presented experimental evaluation nor platform. Any system capable of learning from demonstrations can benefit from this work. Moreover, the framework’s modularity allows loading to its library any primitive skill that might be required for dual-arm manipulation purposes.

Future work will significantly extend the library of primitive skills such that more challenging dual-arm manipulation behaviours can be addressed within the framework. In this regard, imminent efforts will focus on learning force-dependant primitive skills or other actions requiring complex synchronisation between end-effectors, such as the opening of a bottle’s screw cap or succeeding in the peg-in-a-hole tasks.

## ACKNOWLEDGMENTS

This work has been partially supported by ORCA Hub EPSRC (EP/R026173/1) and consortium partners.

## References

1. Ardón, P., Pairet, È., Petrick, R., Ramamoorthy, S., Lohan, K.S.: Reasoning on grasp-action affordances. In: Conference Towards Autonomous Robotic Systems. Springer (2019)
2. Ardón, P., Pairet, È., Ramamoorthy, S., Lohan, K.S.: Towards robust grasps: Using the environment semantics for robotic object affordances. In: Proceedings of the AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy. pp. 5–12. AAAI Press (2018)
3. Argall, B.D., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. *Robotics and autonomous systems* **57**(5), 469–483 (2009)
4. Bajcsy, A., Losey, D.P., O’Malley, M.K., Dragan, A.D.: Learning from physical human corrections, one feature at a time. In: Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. pp. 141–149. ACM (2018)
5. Billard, A., Calinon, S., Dillmann, R., Schaal, S.: Robot programming by demonstration. In: Springer handbook of robotics, pp. 1371–1394. Springer (2008)

6. Gams, A., Ijspeert, A.J., Schaal, S., Lenarčič, J.: On-line learning and modulation of periodic movements with nonlinear dynamical systems. *Autonomous robots* **27**(1), 3–23 (2009)
7. Gams, A., Nemeč, B., Ijspeert, A.J., Ude, A.: Coupling movement primitives: Interaction with the environment and bimanual tasks. *IEEE Transactions on Robotics* **30**(4), 816–830 (2014)
8. Hoffmann, H., Pastor, P., Park, D.H., Schaal, S.: Biologically-inspired dynamical systems for movement generation: automatic real-time goal adaptation and obstacle avoidance. In: *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. pp. 2587–2592. IEEE (2009)
9. Ijspeert, A.J., Nakanishi, J., Hoffmann, H., Pastor, P., Schaal, S.: Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation* **25**(2), 328–373 (2013)
10. Ijspeert, A.J., Nakanishi, J., Schaal, S.: Movement imitation with nonlinear dynamical systems in humanoid robots. In: *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*. vol. 2, pp. 1398–1403. IEEE (2002)
11. Metta, G., Sandini, G., Vernon, D., Natale, L., Nori, F.: The icub humanoid robot: an open platform for research in embodied cognition. In: *Proceedings of the 8th workshop on performance metrics for intelligent systems*. pp. 50–56. ACM (2008)
12. Pairet, È., Ardón, P., Broz, F., Mistry, M., Petillot, Y.: Learning and generalisation of primitives skills towards robust dual-arm manipulation. In: *Proceedings of the AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*. pp. 62–69. AAAI Press (2018)
13. Pairet, È., Hernández, J.D., Lahijanian, M., Carreras, M.: Uncertainty-based online mapping and motion planning for marine robotics guidance. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 2367–2374. IEEE (2018)
14. Park, D.H., Hoffmann, H., Pastor, P., Schaal, S.: Movement reproduction and obstacle avoidance with dynamic movement primitives and potential fields. In: *Humanoid Robots, 2008. Humanoids 2008. 8th IEEE-RAS International Conference on*. pp. 91–98. IEEE (2008)
15. Pastor, P., Hoffmann, H., Asfour, T., Schaal, S.: Learning and generalization of motor skills by learning from demonstration. In: *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. pp. 763–768. IEEE (2009)
16. Pattacini, U., Nori, F., Natale, L., Metta, G., Sandini, G.: An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots. In: *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. pp. 1668–1674. IEEE (2010)
17. Rai, A., Meier, F., Ijspeert, A., Schaal, S.: Learning coupling terms for obstacle avoidance. In: *Humanoid Robots (Humanoids), 2014 14th IEEE-RAS International Conference on*. pp. 512–518. IEEE (2014)
18. Smith, C., Karayiannidis, Y., Nalpantidis, L., Gratal, X., Qi, P., Dimarogonas, D.V., Kragic, D.: Dual arm manipulation: A survey. *Robotics and Autonomous systems* **60**(10), 1340–1353 (2012)
19. Topp, E.A.: Knowledge for synchronized dual-arm robot programming. In: *AAAI Fall Symposium Series 2017*. AAAI Press (2017)
20. Ude, A., Nemeč, B., Petrić, T., Morimoto, J.: Orientation in cartesian space dynamic movement primitives. In: *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. pp. 2997–3004. IEEE (2014)
21. Zöllner, R., Asfour, T., Dillmann, R.: Programming by demonstration: dual-arm manipulation tasks for humanoid robots. In: *IROS*. pp. 479–484 (2004)