Haptic exploration of unknown objects for robust in-hand manipulation

Gokhan Solak, *Member, IEEE*, Lorenzo Jamone, *Member, IEEE*

***Abstract*—Human-like robot hands provide the flexibility to manipulate a variety of objects that are found in unstructured environments. Knowledge of object properties and motion trajectory is required, but often not available in real-world manipulation tasks. Although it is possible to grasp and manipulate unknown objects, an uninformed grasp leads to inferior stability, accuracy, and repeatability of the manipulation. Therefore, a central challenge of in-hand manipulation in unstructured environments is to acquire this information safely and efficiently. We propose an in-hand manipulation framework that does not assume any prior information about the object and the motion, but instead extracts the object properties through a novel haptic exploration procedure and learns the motion from demonstration using dynamical movement primitives. We evaluate our approach by unknown object manipulation experiments using a human-like robot hand. The results show that haptic exploration improves the manipulation robustness and accuracy significantly, compared to the virtual spring framework baseline method that is widely used for grasping unknown objects.**

***Index Terms*—Haptic exploration, Dexterous manipulation, Grasping force optimisation, Learning from demonstration**

I. Introduction

The future of robotics lies in unstructured environments. Robots are expected to operate in environments that are not specifically designed for them, such as daily human environments [1] or even hostile environments [2]. Unlike industrial settings, robots in these contexts need to deal with previously unknown objects and tasks [1], [3]. Coping with such variety of objects and tasks can be facilitated using robot hands that are similar to human hands: in fact, anthropomorphic hands have the obvious advantage of fitting the environments and tools which are designed for humans [4].

Robots operating in an unstructured environment are required to carry out novel tasks on novel objects. We proposed a solution for learning object-centric in-hand manipulation tasks from demonstration [5]. Our solution uses dynamical movement primitives (DMPs) [6] to learn a generalizable and robust representation of the desired motion. The system was able to execute these motions using the virtual spring framework (VSF) [7], an object-agnostic grasping method. However, being object-agnostic comes at the cost of manipulation robustness. The forces that are directed towards the centre of the contact points may generate slips, as shown in Fig. 1. Contact slips perturb the grasp in an uncontrolled way, reducing the in-hand stability of the object.

G. Solak and L. Jamone are with ARQ (Advanced Robotics at Queen Mary), School of Electronic Engineering and Computer Science & School of Engineering and Materials Science, Queen Mary University of London, London, UK.

For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

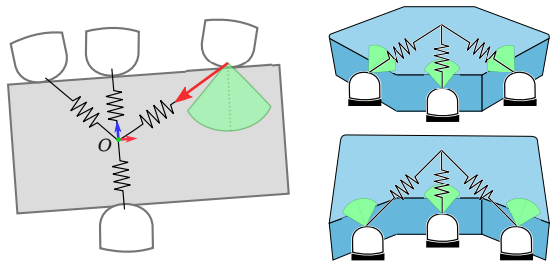

(a) (b)

Fig. 1. (a) When the robot fingertips are connected to the object frame O with virtual springs, the spring force (red) may become tangential to the object surface, fall outside the friction cone (green) and cause a contact slip. Ideally, all grasping forces must remain in the friction cone. (b) An illustration of the classical virtual springs and friction cones (green) on different object shapes. Below, finger slips occur as the spring forces are outside the friction cones, while above the springs do not cause a slip.

In this work, we propose a novel haptic exploration method for extracting necessary object properties to improve the in-hand manipulation robustness and accuracy [8], [9]. Some of the object properties (shape, pose, colour) can be obtained by vision; however, some (friction, centre of mass, inertia parameters) require tactile interaction with the object. For shape and pose, touch sensing is still useful in scenarios where vision modalities are not available or the object is occluded [8]. In this work, we use only the touch interaction to extract the surface normal and friction information at contact points. These properties are chosen as they are required to calculate and satisfy the friction constraints and prevent unwanted slips.

Our exploration method is designed to be *safe* and *efficient*, which are the requirements of operating in an unstructured environment. In this context, *safety* is related to the manipulated object. Losing or damaging the object at hand means a failure of the task, and it may have further consequences depending on the application. *Efficiency* is the requirement that the exploration concludes in a short time without interruption.

When an object shape is fully known, the motion can be constrained to prevent unsafe interactions [10]. However, the shape assumption is not realistic in unstructured environments. The safety problem is often avoided by fixing the object pose physically [11], [12], however, methods that work in the wild cannot impose such constraints. Some methods depend on unsafe actions such as gravity-induced object slips to observe the object properties [13], [14]. We apply the exploratory actions while the object is lying on a supporting plane to avoid dropping it unsafely. Our multi-fingered setup allows inducing slips without depending on external forces, such as gravity-induced forces. As a consequence, we continuously balance the grasping forces during exploration so that the object is not moved or lost (Sec. IV-B4).

With the combination of haptic exploration and learning from demonstration, an in-hand manipulation framework for unknown objects is formed. This framework aims to obtain the two components of missing information in unstructured environments. Motion information is learned from demonstration of a human user and object information is extracted by autonomous haptic exploration as illustrated in Fig. 2. We use the extracted object information to define a grasping force optimisation (GFO) problem which is solved to determine the fingertip forces that satisfy the friction constraints (Sec. IV-D).

We evaluate the effect of haptic exploration directly on the performance of in-hand manipulation tasks on unknown objects. We execute the tasks that are learned from demonstration using a four-finger anthropomorphic robot hand, on objects of different shapes and surface materials. We compare our exploration-based approach to our baseline, the VSF, which is widely used for object-agnostic grasp and manipulation [7], [15]–[18], including our previous work [5]. Our current work differs from [5] as we add a haptic exploration step before the manipulation and apply GFO to maintain a better grasp. The results show that haptic exploration improves the manipulation robustness and accuracy significantly.

Our contributions are summarised as follows:

- We propose a novel haptic exploration method that caters to work safely and efficiently in real-world scenarios, that is, without damaging or losing the object, and concluding in a short time.
- We evaluate the effect of haptic exploration directly on in-hand manipulation of unknown objects, through real-world robot experiments. We publish the data collected in these experiments [19].

We begin with the discussion of the related work in Sec. II. Then we give the overview of our framework in Sec. III and the details of its modules in Sec. IV. Lastly, in Sec. V & VI we report our experiments and conclusions.

II. Related work

A. Haptic exploration

Tactile perception is a vital skill for a wide range of robotic applications such as grasping, locomotion, human-robot interaction and manipulation [9]. Tactile perception methods can be categorised by being online or offline, data-driven or analytic, and by the object properties they target [8], [9].

Tactile perception may happen during task execution [13], [14], [20] or in a dedicated exploration phase [21]–[23]. As an exception [24] integrates the exploration strategy into task execution to discover the object shape while rotating it. We use an exploration strategy because we try to gain object properties information and the exploration requires visiting possibly unstable states that are not safe for task execution.

The exploration approaches can also be categorised as data-driven or analytical [9]. Our approach is analytical which has the advantages of not depending on data collection and being human-interpretable. However, data-driven approaches enable working with the elements for which precise modelling is difficult such as soft robots [25], deformable objects [26], and complex sensor inputs [27]. As shown recently [28], they can also be used for decision-making during the exploration process.

Haptic exploration applications target different object properties [8] such as: the global object shape [21], [29]; the local object geometry [30], [31]; the surface texture [11], [12]; the stiffness [14], [22]; the centre of mass [10], [22]; or the friction coefficient [32], [33]. In this work, we aim to extract the local surface normals and the friction coefficient information as required by our grasp optimisation approach. Hence, we discuss the works on the friction property more in detail below.

The friction properties of an object are mainly useful in avoiding contact slips and improving grasp stability [32]. Therefore, some approaches rely on online slip detection, without estimating the friction, to apply reactive control actions when an incipient slip is detected [14], [20]. The online slip detection approaches can be built on data-driven models in systems that are hard to model, such as soft hands [25]. Other approaches estimate the static friction coefficient μ_s which indicates the shear/normal force ratio when the slip occurs. Estimation of μ_s is done either on the initial contact or by letting a *gross slip* occur [32]. The former is achieved using specially designed tactile sensors that measure strain [34] or make multiple contacts with the object at different angles [35]. Another specific sensor encourages *incipient slips* to measure μ_s before a gross slip occurs [36]. However, having an ad-hoc friction sensor is costly and the presented sensors are large in size. The methods that let a gross slip happen usually use gravity to cause slips while lifting the object [13], [14]. This approach carries a risk of dropping the object, which is undesired when the object is fragile or it contains liquids. Also, it is harder to induce gravity-based slips when the object is lightweight. If the exact object shape is known, the fingers can be commanded to slip tangentially to the surface while measuring the friction coefficient [10]. However, we do not assume a known object model. Some works fix the object pose physically [11], [12], however, this is not possible in autonomous scenarios.

Our method depends on slip detection to identify the friction cone at a contact point. Without prior knowledge of the friction model, slips can be detected by analyzing the vibration frequencies [37], [38], comparing the normal and tangential force ratio [13], detecting contact displacements with a camera-based sensor [39], modelling the uncertainty in time-series data [40] or training a classifier on tactile arrays [20], [41]. Vibration-based approaches can be applied with affordable sensors, however, they may suffer from the vibration noise from other sources [33], [42]. Initially, we tried to use a vibration-based approach for slip detection. However, it generated a high false-positive rate due to the noise. In the current version, we are using a position-based approach for slip detection, but it could be replaced with a tactile-based approach if better sensors were available, e.g. more sensitive [43], [44] and distributed over a larger surface [45], [46].

In summary, our haptic exploration method is a novel way of estimating friction, using simple sensing modalities, without object shape knowledge and without relying on gravity. Our experiments are novel in the sense that we show the effect of haptic exploration directly on in-hand object manipulation

performance, while the existing works focus on grasp stability without motion.

B. In-hand manipulation

In-hand manipulation (or in earlier literature, dexterous manipulation) is the problem of manipulating an object w.r.t. the hand frame [47]. The methods dealing with this problem have to determine the actuator forces/torques that achieve the desired object motion [48]. In this work, we focus on manipulation using high-DoF multi-fingered robot hands, as opposed to under-actuated grippers that exploit the environment constraints [49] or specific hand dynamics [50] to achieve in-hand manipulation.

The in-hand manipulation problem can be divided into motion planning, which determines the desired trajectory of the object; and control, which executes the manipulation safely and accurately [51]. For motion planning, we use the learning from demonstration (LfD) approach that we proposed earlier [5], for its data efficiency and generalisation capabilities.

An in-hand manipulation controller must deal with both keeping the object in grasp and performing the object motion simultaneously [51], unlike the classical manipulation problem where the object is held tightly by a single manipulator. Answering these requirements is possible by careful modelling of the robot, object, and motion [52]. However, such information is often not available in real-world manipulation scenarios [1].

One way to deal with the missing object knowledge is increasing sensing capabilities, such as tactile sensing [53] or in-hand object tracking [54]. However, when such sensing is not available, it is possible to mitigate the uncertainties in object information by adding active or passive compliance to the system. Recent works emphasise the importance of physical compliance by achieving complex in-hand manipulation using soft hands [55], [56]. Where such hardware is not available, a common way to manipulate unknown objects is the VSF, that is, creating a virtual coupling between the robot fingers to apply the grasping forces [7], [16]–[18]. In grasping, it is shown to increase robustness under uncertainty [18]. A common aspect of these works is to define the object frame w.r.t. the robot fingertips [57]. [7], [18] assign a constant stiffness for the virtual springs. However, [16] learns the impedance parameters from demonstration and [17] learns a probabilistic model of stable grasps which allows online adaptation of spring stiffness and rest length parameters. We also adopted the VSF in our previous work [5]. Another approach to active compliance regulates the common grasp stiffness instead of individual fingers to achieve a task-dependent stiffness ellipsoid [58]. In the past VSF works, the grasping forces are directed towards the virtual object frame, which may cause large contact slips with non-round objects as explained in Fig. 1. Our work differs as we extract object information by haptic exploration to optimise the grasp.

Given an already stable grasp, non-compliant in-hand manipulation control can also be formulated without an object model. [59] excludes all sensing except the onboard sensors to achieve stable *blind-grasping*, however, it requires the friction coefficient and local surface normals to be available. Our haptic exploration method could be integrated with this controller to provide this information. [60] proposes a purely kinematic trajectory optimiser that minimises the displacement of the contact positions from the initial grasping points, rather than keeping fixed contact points, for improved manipulation accuracy. This could be an alternative over our no-slip assumption in manipulation control too. However, the difference of our method is that the motion plan is learned from demonstration, for the added value of the trajectory shape, while in [60] only the final pose matters. Additionally, we report the contact stability using fingertip force sensing in our experiments.

In another related work [54], online optimisation adapts the grasping forces, instead of directing them towards the object centre. The surface friction coefficient is assumed to be known, and the surface normals needed for the optimisation are estimated using vision-based in-hand object localisation. The main differences with our work are that: we do not assume friction to be known; we use haptic exploration (instead of vision) to estimate the normals.

Notably, some recent reinforcement learning methods use a single model to answer both motion planning and low-level control problems [61]–[63]. These methods learn highly dexterous behaviour through trial and error. However, they were only applied to manipulation scenarios that facilitate safe repetition [62], [64] since they require long training times.

For control, we use a hybrid force/position controller in which the internal forces are determined by the GFO. GFO aims to balance the internal forces acting on the object while satisfying the grasping constraints. Definition of the grasping constraints requires the local surface normals and friction information which we obtain through haptic exploration. The expression of the optimisation problem is important for the computation complexity and solution quality [65]. The source of complexity in these problems is the non-linearity of the objective and constraint functions. For this reason, early works solved a linear approximation of the problem using linear programming [66]. Different formulations such as convex optimisation or second-order cone problems were used alongside specific algorithms to increase the efficiency [67]. A comparison shows the accuracy advantage of the non-linear optimisation approach with a speed trade-off against the linearised approaches [65]. We adopt a non-linear approach in this paper, as our experiments show that it is sufficiently fast for our requirements.

In this work, we focus on manipulations within the workspace of the fingers (*coordinated manipulation* [47]) and proceed to evaluate our haptic exploration contribution on in-hand manipulation. Answering the other types of dexterous manipulation problems where the contact points relocate (*finger gaiting* [47]) requires object shape information, of which our method remains uninformed.

III. Framework

In this section, we describe the overall structure of our in-hand manipulation of unknown objects framework, and dependencies between the modules. The details of the modules are described later in Sec. IV. We can examine the framework in two sections as depicted in Fig. 2: Knowledge acquisition

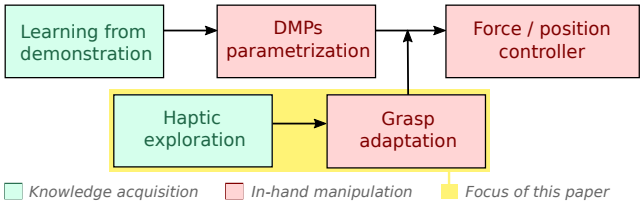

Fig. 2. Our in-hand manipulation framework aims to acquire the missing information of both the motion and the object efficiently before manipulation. The motion is learned from demonstration of a human user and the object information is extracted by autonomous haptic exploration. This information is used during manipulation to achieve the task robustly. The object information is used to adapt the grasp during manipulation and the learned movement primitives are parametrised for the specific task goals. The manipulation is achieved by a hybrid force/position controller. This paper focuses on the haptic exploration and grasp adaptation components, highlighted by yellow.

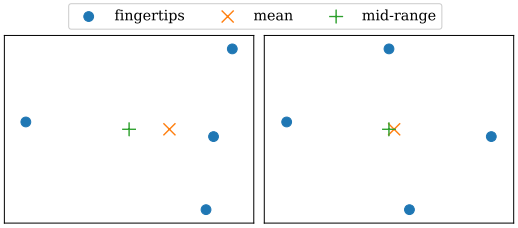

Fig. 3. Virtual object position calculated using different centre measures with asymmetrical (left) and symmetrical (right) hand designs. When the fingertips are asymmetrically distributed, the *mean* centre is biased towards the side that has more fingers. We want the object position to be close to the geometrical centre of the object, regardless of the hand design. As seen in the figures, the *mid-range* measure gives a good approximation in both designs while the *mean* is biased by finger density.

(Sec. III-A), which happens before manipulation; and in-hand manipulation (Sec. III-B), which achieves control during manipulation.

A. Knowledge acquisition

Our system includes two modes of knowledge acquisition: extracting the object information by haptic exploration (Sec. IV-B) and learning the task by LfD (Sec. IV-C). The knowledge of the task and of the object are then used by the in-hand manipulation controller to realize the task robustly with a given object. This paper focuses on the haptic exploration component. The LfD component is presented in a previous work [5].

The two components are separate from each other, and do not happen simultaneously. The LfD component is used when a new task has to be learned; the haptic exploration is used when a known task is applied to a new object. Both components acquire the required information efficiently, i.e. without requiring long training times: LfD achieves this using DMPs, which learn a motion primitive from a single demonstration [6], and the haptic exploration achieves this by a short interaction procedure (described extensively in this paper).

Our implementation of the haptic exploration does not require any tactile sensing; however, additional tactile sensing could improve the slip detection, which is part of our haptic exploration, and the opposing thumb assumption that we discuss in Sec. IV-B3 might be relaxed.

B. In-hand manipulation

Our in-hand manipulation method combines a force controller for grasping and a position controller for manipulation (Sec. IV-E). The force controller is modulated by the desired forces that are calculated by the GFO (Sec. IV-D). The GFO needs the surface normal and friction information that are extracted by haptic exploration (Sec. IV-B). On the other hand, the position controller needs the desired object trajectory, which is obtained by parametrisation of the DMPs which are learned from demonstration (Sec. IV-C).

IV. Methodology

In this section, we present the details of the modules of our framework. We begin with the definition of the virtual object frame that is central to our method (Sec. IV-A). Then we describe the modules of the knowledge acquisition system: haptic exploration (Sec. IV-B) and LfD (Sec. IV-C); and the in-hand manipulation modules: the GFO (Sec. IV-D) and the manipulation controller (Sec. IV-E). Lastly, we describe our baseline method in Sec. IV-F.

A. Virtual Object Frame

We use a virtual object frame as the frame of manipulation as we assume no object information. We use the virtual object frame for multiple purposes in our framework. Firstly, we use it as the object pose when recording and reproducing DMPs. Secondly, we use it to track the expected contact points during the exploration so that a slip can be distinguished from object movement (Sec. IV-B3). Lastly, we use it during manipulation to track the object pose and the relative contact poses, so that the forces can be calculated in the object frame of reference.

The virtual frame is calculated only using the fingertip positions $\mathbf{p}_{ci}$, without any visual information. We follow the existing literature in estimating the position and the rotation of the object frame [57]. The virtual object frame $\mathbf{x_o} = (\mathbf{p_o}, \mathbf{R_o})$ denotes the approximate pose of the object:

$$\mathbf{p}_o = \frac{1}{n}\sum_{i=1}^{n} \mathbf{p}_{ci} \tag{1a}$$

$$\mathbf{R}_o = [\hat{\mathbf{r}}_x, \hat{\mathbf{r}}_y, \hat{\mathbf{r}}_z] \in SO(3) \tag{1b}$$

$$\mathbf{r}_x = \frac{\mathbf{p}_{c1}-\mathbf{p}_{c3}}{\|\mathbf{p}_{c1}-\mathbf{p}_{c3}\|} + \frac{\mathbf{p}_{c2}-\mathbf{p}_{c4}}{\|\mathbf{p}_{c2}-\mathbf{p}_{c4}\|}$$

$$\mathbf{r}_z = (\mathbf{p}_{c1} - \mathbf{p}_{c3}) \times (\mathbf{p}_{c1} - \mathbf{p}_{c3})$$

$$\mathbf{r}_y = \mathbf{r}_z \times \mathbf{r}_x$$

In the calculation of the frame orientation $\mathbf{R}_o$, each axis is normalised as $\hat{\mathbf{v}} = \mathbf{v}/\|\mathbf{v}\|$. We use the hat operator to indicate unit vectors and double vertical columns $\|.\|$ to indicate the L^2-norm in this paper. n is the number of contact points.

This definition of frame orientation (1b) is specific to 4 contact points, however, it is possible to derive it for different number of contacts.

In this work, we modify the position equation (1a) to use the *mid-range* instead of the *mean* centre of the fingertips. Because, for asymmetrical hand designs the *mean* centre is biased towards the side that contains more fingers, as explained in Fig. 3. In the case of the human-inspired Allegro Hand, one thumb opposes three fingers, hence the mean is closer to the fingers, instead of being close to the geometrical centre of

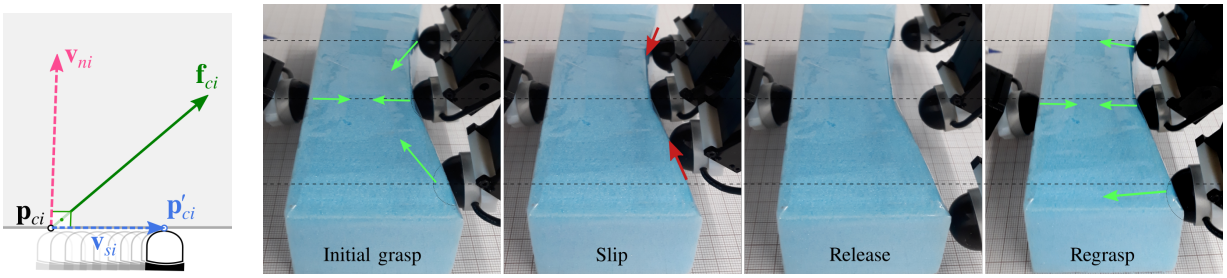

(a) (b)

Fig. 4. (a) In *stage 1*, the object is grasped initially with forces towards the centre and the contact shift vector $\mathbf{v}_{si}$ is observed. Using this information, the approximate surface normal $\mathbf{v}_{ni}$ is calculated. (b) Snapshots of *stage 1*, on *object 3*. Green arrows illustrate the applied forces and red arrows illustrate the slips. The robot applies a central grasp and observes the slip. Then, the object is regrasped with grasping forces along the normals, causing no slip. The initial object pose and contact locations are preserved.

Algorithm 1 *Stage 1* procedure

Initial grasp at points $\mathbf{p}_{ci}$
Set $\mathbf{f}_{ci}$ towards $\mathbf{p}_o$
$\text{slip}_i := 0$
$t :=$ time in *seconds*
repeat
    Apply balanced grasping forces (6)
    $\text{slip}_i :=$ Check for slip (5)
until $\text{slip}_i = 1$ **or** time() $> t + 3$ ▷ slip or $3s$ passed
if $\text{slip}_i = 1$ **then**
    Estimate contact normal $\mathbf{v}_{ni}$ (2)
else
    $\mathbf{v}_{ni} \leftarrow \mathbf{f}_{ci}$

the object. The *mid-range* approach works reliably with both symmetrical and asymmetrical designs as seen in Fig. 3.

The mid-range centre is calculated as follows: $\mathbf{p}_{o,j} = (\max_i \mathbf{p}_{ci,j} + \min_i \mathbf{p}_{ci,j})/2$. Each entry of the centre point $\mathbf{p}_o$ is calculated as the average of the maximum and the minimum values of its dimension over all contact points $\mathbf{p}_{ci}$, eliminating n, the number of fingers; where $\mathbf{p}_o = (\mathbf{p}_{o,1}, \mathbf{p}_{o,2}, \mathbf{p}_{o,3}) \in \mathbb{R}^3, \mathbf{p}_{ci} = (\mathbf{p}_{ci,1}, \mathbf{p}_{ci,2}, \mathbf{p}_{ci,3}) \in \mathbb{R}^3$.

B. Haptic exploration

We use haptic exploration to gain useful information about the object to enhance the grasp robustness. The exploration is done after the initial grasp, before lifting the object for manipulation. The objective of the exploration is to estimate the local surface normals and friction cones. This information is then used in the GFO to set the constraints.

We propose an exploration procedure that consists of two stages. The exploration begins with *Stage 1* when the initial contact with the object is made. *Stage 1* uses the information of the initial grasp to estimate the surface normals $\mathbf{v}_{ni}$ (Alg. 1). Then, the object is re-grasped at the same pose, applying forces alongside the surface normals, and the next stage is executed. *Stage 2* performs a more detailed search to estimate the friction cone which is characterised by the tangent plane orientation and the friction coefficient between contacting materials (Alg. 2). We present the detailed descriptions of the exploration stages in the following subsections.

1) Stage 1: Given a set of grasping points on the object, without knowing the local surface normals and the friction coefficient, a common way to grasp the object is by applying forces towards the geometric centre ($\mathbf{p}_o$) of the grasping points. This grasp may cause contact slips if the object does not have a round shape. Thus, we use this approach only for the initial grasp and we observe whether a slip occurs (Alg. 1). If the initial grasp does not cause a slip, then the object can be handled without any exploration. If a slip occurs, the slip vector $\mathbf{v}_{si}$ lies approximately on the plane that is tangential to the object surface at the initial contact point as shown in Fig. 4.a. We calculate the normal vector $\mathbf{v}_{ni}$ using this information as follows:

$$\begin{aligned} \mathbf{v}_{si} &= \mathbf{p}'_{ci} - \mathbf{p}_{ci} \\ \mathbf{v}_{f||s} &= (\mathbf{f}_{ci} \cdot \hat{\mathbf{v}}_{si})\hat{\mathbf{v}}_{si} \\ \mathbf{v}_{ni} &= \mathbf{f}_{ci} - \mathbf{v}_{f||s} \end{aligned} \tag{2}$$

In summary, $\mathbf{v}_{ni}$ is the vector that is perpendicular to $\mathbf{v}_{si}$ and coplanar with $\mathbf{f}_{ci}$. After calculating the surface normal, the object is regrasped to restore the initial grasp point $\mathbf{p}_{ci}$. Snapshots of a *stage 1* exploration is shown in Fig. 4.b.

2) Stage 2: In this stage, our aim is to determine the friction cone at the contact. We assume that the object has a uniform surface texture, hence the friction coefficient is the same for all contacts.

Because the friction cone is normal to the surface and symmetrical about the normal, it is sufficient to explore one section of the friction cone. We explore the section that contains both the surface normal and the initial grasping force vector as shown in Fig. 5. This way, the explored forces stay between the initial force and the surface normal of that contact.

The exploration procedure consists of regrasping the object multiple times, applying a different grasping force at each iteration. Depending on whether or not a slip occurs with that force, we update the *minimum slipping angle* α_μ, that is, the smallest force/normal angle that caused a slip. α_μ also defines the edge of the approximate friction cone.

A move step rotates the contact force about an axis $\mathbf{v}_R$ that is orthogonal to the plane of search: $\mathbf{v}_R = \mathbf{v}_{ni} \times \mathbf{f}_{ci}$. The next force to explore $\mathbf{f}'_{ci}$ is generated as:

$$\mathbf{f}'_{ci} = \mathbf{R}(\mathbf{v}_R, \theta) \times \mathbf{f}_{ci}. \tag{3}$$

$\mathbf{R}(\mathbf{v}_R, \theta)$ is a rotation matrix, equivalent to a rotation of angle θ about the axis $\mathbf{v}_R$. The choice of θ is crucial for the speed of the exploration. Scanning an arc of $60°$ linearly in steps of

1° requires 30 steps on average. This would be impractical in means of time. Therefore, we follow a binary search approach to explore the cone section.

The binary search observes the middle point of the unexplored space at each iteration. Depending on the outcome of the grasp, it chooses one of the two half-spaces to continue, as illustrated in Fig. 5. Let θ_t be the rotation angle at step t. In the beginning, θ_0 is the angle between the centre-bound force $\mathbf{f}_{ci}$ and the normal $\mathbf{v}_{ni}$. Given that the finger slipped at the initial grasp, we rotate halfway towards the normal: $\theta_1 = -\theta_0/2$. At the next step, the rotation angle is halved again, however, the sign of the rotation depends on whether a slip occurs:

$$\theta_{t+1} = \begin{cases} -|\theta_t|/2 & \text{, slip} \\ |\theta_t|/2 & \text{, otherwise} \end{cases} \tag{4}$$

The move step is repeated until the desired accuracy is established. We terminate the search after 5 steps. Scanning an arc of 60° in 5 steps decreases the unexplored space to $60/2^5 = 1.875^\circ$. We find it sufficient in this study. In the end, the friction coefficient is determined as: $\mu = \arctan(\alpha_\mu)$.

The *stage 2* procedure is applied only on the fingers that slip at the initial grasp. The absence of slip indicates that the initial force is good for this contact, hence no improvement is required. In the example of Fig. 4.b, the thumb and the middle finger do not slip, thus they are not involved in the rest of the exploration. They maintain a force towards the centre during exploration and manipulation.

The rotation of contact forces during the exploration may break the constraint of zero net force and move the object. This is undesired as it may lead to the loss of the object or misleading slips during the search. To avoid this, the forces are balanced during the exploration as described in Sec. IV-B4. The balancing optimises the magnitudes of the forces in order that the total wrench acting on the object is minimised.

An important building block of our haptic exploration method is slip detection. It is used to direct the search and mark the cases where the friction cone is violated.

3) Slip Detection: In our work, we use fingertip displacement as the indicator of gross slips. Gross slips create observable changes in the contact positions, hence after each search step, the object must be regrasped.

A slip is detected when the distance between the contact positions before ($\mathbf{p}_{ci}$) and after ($\mathbf{p}'_{ci}$) the grasp exceeds a threshold a_{slip} (H is the Heaviside function):

$$\text{slip}_i = H(\|\mathbf{p}'_{ci} - \mathbf{p}_{ci}\| > a_{\text{slip}}). \tag{5}$$

We determined a_{slip} as 1.3 *mm* considering the encoder sensitivity and slip performance in our setup.

This approach to slip detection has the advantage of not requiring a tactile sensor, however, it is sensitive to the contact positions. Although we balance the forces as described in the next subsection, the object can move and rotate slightly during the exploration. When detecting slips, we must disregard the displacements due to the object movement, i.e., when all contacts move together without slipping.

We use the virtual frame estimation as described in Sec. IV-A to calculate the expected position of each contact

Algorithm 2 *Stage 2* procedure

$\theta_i \leftarrow \mathbf{f}_{ci}\angle\mathbf{v}_{ni}$ ▷ $\mathbf{v}_1\angle\mathbf{v}_2$: angle between $\mathbf{v}_1$ and $\mathbf{v}_2$
loop(5 times)
  Regrasp at $\mathbf{p}_{ci}$
  Apply balanced grasping forces (6)
  $\text{slip}_i :=$ Check for slip (5)
  if $\text{slip}_i = 1$ **then**
    $\alpha_\mu \leftarrow \mathbf{f}_{ci}\angle\mathbf{v}_{ni}$
  Update θ_i (4)
  Rotate $\mathbf{f}_{ci}$ by θ_i (3)
Calculate $\mu = \arctan(\alpha_\mu)$

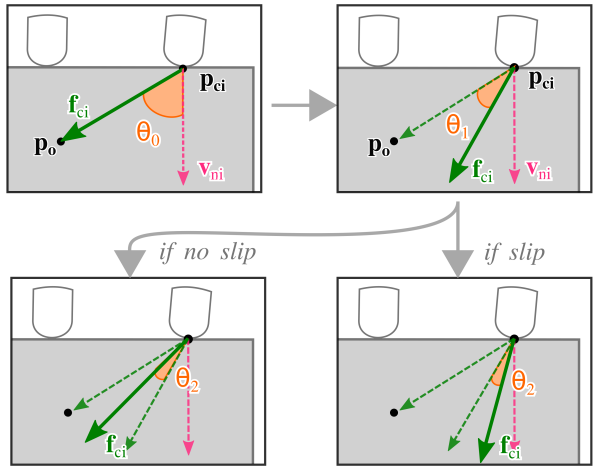

Fig. 5. *Stage 2* searches the local space of forces until finding the minimum slipping angle. The force takes different orientations between the initial force $\mathbf{f}_{ci}$ and the estimated normal $\mathbf{v}_{ni}$ on the plane that contains these two vectors. The force is rotated to divide the space between the minimum slipping angle and the maximum no-slip angle in half, as in a binary search. The next half-space to search is determined by the slip status at the current angle.

point, i.e. the contact position if the initial grasp is preserved (no slip occurred). However, the position of the virtual frame shifts when a contact point slips, creating the delusion that the object moved. For this reason, instead of using the virtual object position, we use the thumb position as the reference point. Since our robot hand has an opposing thumb design, the thumb is unlikely to slip without moving the object. Following this observation, we calculate the expected contact points by their relative position to the thumb. Thus, slip detection without tactile sensing requires the assumption of a robot hand with an opposing thumb design. This assumption could be relaxed using a more advanced tactile sensor for detecting slips [33].

4) Balancing forces: We model this step as a bound-constrained linear least-squares optimisation problem, to find the coefficients c_i that scale the unit grasping forces $\hat{\mathbf{f}}_{ci}$, so that, the total object wrench is minimised. The coefficients are bounded between the lower and upper bounds (c_{low}, c_{up}):

$$\underset{c_i}{\text{minimise}} \quad \frac{1}{2}\|\mathbf{0} - \mathbf{G}\mathbf{h}_c\|^2 \tag{6a}$$

$$\text{subject to} \quad c_{low} \leq c_i \leq c_{up} \tag{6b}$$

$\mathbf{G}$ is the grasp matrix that maps the contact forces ($\mathbf{h}_c = [\mathbf{f}_{c1}\mathbf{f}_{c2}\ldots\mathbf{f}_{cn}]^\mathsf{T}$) to the object wrench. Each contact force

is determined as the unit contact force scaled by its optimal coefficient: $\mathbf{f}_{ci} = c_i \hat{\mathbf{f}}_{ci}$.

Please note, this is an underdetermined problem, that does not always have an optimal solution. As a result, the object wrench may become non-zero during the exploration, but it is still minimised as much as possible. The supplemental material shows the outcomes of force balancing during exploration.

The lower bound must be greater than zero ($c_{low} > 0$) to avoid the trivial solution of all zeros. We also choose an upper bound for safety.

Compared with the GFO during the manipulation (Sec. IV-D), this optimisation problem does not change the directions of the forces, only adjusts their magnitudes. Therefore, it has a lower dimensionality.

C. Learning from demonstration

We follow the LfD approach that is proposed in previous work [5]. This method uses DMPs [6] to learn an in-hand object motion from demonstration. The learning method is object-centric, i.e., the learned DMPs encode the object trajectory. The movement primitives can be parametrised to meet the task requirements without losing the trajectory shape.

Task demonstration is delivered kinesthetically, by physically holding the robot and moving it to perform the desired task. The controller applies grasping forces via the VSF and compensates gravity during demonstration, however, the position control is deactivated. The robot records the motion of its fingertips and calculates the virtual object frame as explained in Sec. IV-A.

The recorded object trajectory is used to learn a set of DMPs. Each DMP describes an independent dimension of the object pose. We store 3 DMPs for the position and 6 DMPs for the first two unit axes of the rotation matrix. The third axis is calculated later as a cross product of the first two. Please note, since the rotation matrix representation is redundant, i.e., members of the rotation matrix are not independent, the learned DMPs may break the orthogonality of the axes. Better treatment of Cartesian space orientation with DMPs is proposed in [68].

When the task reproduction is requested, the learned DMPs are modulated with the user-specified task conditions: the initial object pose, the duration, and the final object pose. These parameters are enough to create an object trajectory that retains the shape of the learned motion. The object trajectory is executed by the position controller as described in Sec. IV-E. The details of the DMPs are omitted in this paper as we follow the same DMP formulation as in the previous work [5].

D. Grasping Force Optimisation

We solve a nonlinearly constrained optimisation problem to find the set of minimal contact forces $\mathbf{h}_c$ that satisfies the constraints of a robust grasp:

$$\underset{\mathbf{h}_c}{\text{minimise}} \quad \sum_{i=1}^{n} \|\mathbf{f}_{ci}\| \tag{7a}$$

$$\text{subject to} \quad \mathbf{G}\mathbf{h}_c = \mathbf{0} \tag{7b}$$

$$\|\mathbf{f}_{ci}^t\| / \|\mathbf{f}_{ci}^n\| < \mu \tag{7c}$$

$$b_{low} \leq \|\mathbf{f}_{ci}\| \leq b_{up} \tag{7d}$$

We impose 3 types of constraints:

- *Zero object wrench*: The net wrench applied on the object frame is null, so that the object pose is not perturbed by the grasp controller (7b).
- *Friction cone*: Each contact force $\mathbf{f}_{ci}$ satisfies the Coulomb friction equation where μ is the friction coefficient, $\mathbf{f}_{ci}^n$ and $\mathbf{f}_{ci}^t$ are the normal and tangential components of the contact force (7c). These components are calculated using the surface normal vector $\hat{\mathbf{v}}_{ni}$ at that contact point.
- *Force limits*: We enforce lower and higher bounds on the contact force magnitudes (7d). The lower boundary b_{low} is to apply sufficient force to lift the object, the upper boundary b_{up} is to keep the robot and the object safe.

This problem requires the friction coefficient μ, the normal vectors $\hat{\mathbf{v}}_{ni}$ and the boundary constraints (b_{low}, b_{up}) to be specified. Boundary constraints are set by the user, and the normals and the friction coefficient are estimated by the haptic exploration method (Sec. IV-B). We assumed hard point contacts with friction (PCWF) when forming the grasp matrix. In comparison to balancing the forces during haptic exploration (Sec. IV-B4), the force directions are also variable in this problem. Hence it is more flexible, but also more complex.

E. Manipulation controller

The controller combines the force control signal $\mathbf{u}_f$ for grasping, the position control signal $\mathbf{u}_p$ for manipulation, and the gravity compensation signal $\mathbf{u}_g$:

$$\mathbf{u}_c = \mathbf{u}_f + \mathbf{u}_p + \mathbf{u}_g \tag{8}$$

$\mathbf{u}_c, \mathbf{u}_f, \mathbf{u}_p, \mathbf{u}_g \in \mathbb{R}^m$ are vectors of torque commands for the dexterous hand actuators. m is the total degrees of freedom of the robot hand. $\mathbf{u}_g$ is the set of actuator torques that cancel the gravity effect on the robot fingers without any object. Note the potential conflict between the force ($\mathbf{u}_f$) and the position ($\mathbf{u}_p$) control signals, i.e., the grasping force signal may perturb the object motion. For this reason, the GFO constrains the grasping forces to induce zero wrench on the object, as expressed in (7a).

The grasp torques are computed as $\mathbf{u}_f = \mathbf{J}_h^\intercal \mathbf{h}_c$ using the desired contact forces $\mathbf{h}_c$ from the GFO (Sec. IV-D), and the hand Jacobian $\mathbf{J}_h$ [47].

We use PID control to produce $\mathbf{u}_p$ to follow a desired object trajectory that is produced using the learned DMPs. We first compute the fingertip trajectories by applying the current relative transform from the object frame to the fingertips on the whole object trajectory, adhering to our *no slipping and rolling assumption*. Then, we apply inverse kinematics to create a joint trajectory for each finger. These joint trajectories are used as targets of the joint position controller.

We compare our haptic exploration method to the baseline of the VSF. The baseline manipulation controller is the same except that the grasping forces are generated by the VSF instead of the GFO, without any object knowledge.

F. Baseline

The VSF generates the desired forces using virtual springs that apply force towards the object frame. The force applied

by each spring is calculated individually using Hooke's law. The grasping force for finger i:

$$\mathbf{f}_{ci} = K_{ci}(\|\mathbf{v}_{fi}\| - L_i) \tag{9}$$

$\mathbf{v}_{fi} = \mathbf{p}_o - \mathbf{p}_{ci}$ is the vector between the i^{th} contact point $\mathbf{p}_{ci}$ and the virtual object position $\mathbf{p}_o$. L_i is the rest length of the spring connecting i^{th} fingertip to the virtual frame origin. We initialise the rest length proportionally to the measured distance of the contact to the object centre when the fingers are in touch with the object ($L_i = 0.9\,\|\mathbf{v}_{fi}\|$). As a result, the springs start applying the base grasping forces.

K_{ci} is the stiffness of the i^{th} spring determining the strength of the coupling between the ends of the spring. We assign a hand-tuned value to K_{ci} depending on the object weight.

V. Experiments

We evaluate our contribution by teaching the robot hand different tasks by demonstration and reproducing these tasks on different objects. We evaluate the effect of haptic exploration by comparing it to the baseline grasp with the virtual springs. We run in-hand manipulation tasks with both approaches and compare their performances. The results show that the haptic exploration procedure improves contact stability and manipulation accuracy significantly.

We first define our problem and performance measures in Sec. V-A. Then, we present our setup (Sec. V-B) and dataset (Sec. V-C). In Sec. V-D, we give the details of the grasp and manipulation sequence that is used to collect the data. The results are presented in Sec. V-E and discussed in Sec. V-F.

A. In-hand manipulation problem

The problem we focus on is *robust* in-hand manipulation of unknown objects. *In-hand* implies that the object pose is transformed w.r.t. the hand frame. Thus, the robot arm that carries the hand does not move during manipulations. We measure *robustness* by two criteria: stability of the finger contacts on the object, namely the *contact stability*; similarity of the executed motion to the learned motion, namely the *accuracy*.

The haptic exploration works in the literature often use a grasp quality metric to show improvement, however, in our work the grasp does not change kinematically. Thus, we report the *contact stability* as the direct outcome of more stable grasps. This metric is also affected by the dynamics of in-hand manipulation, thus highlighting the importance of friction-aware force modulation. On the other hand, related works generally present accuracy of the final object pose, however, we report the *accuracy* through the entire trajectory as our goal is to reproduce the desired motion as a whole.

1) Contact stability: We evaluate the contact stability during manipulation by the persistence of the touch signal on the fingertips. The absence of touch means that contact with the object is temporarily lost. In a robust grasp, the contacts should be maintained throughout the manipulation. This metric is quantified as the ratio of time in touch to the total time T, for all n fingers:

$$\sum_{i=0}^{n}\sum_{t=0}^{T} H(\|\mathbf{f}_{si}(t)\| - a_f)/nT \tag{10}$$

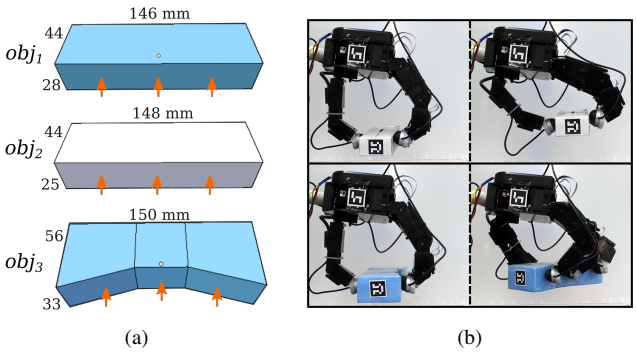

(a) (b)

Fig. 6. (a) Illustration of the objects that are used in the experiments, their dimension sizes (mm) and expected contact locations (orange). The size differences are the result of manual object production process. (b) We apply the learned object motions on different objects to validate the method. In *task 1* (above), the object follows an arc in the right-up direction, and in *task 2* (below), the object is pulled upwards and rotated.

$\mathbf{f}_{si}(t)$ is the measured force for finger i at time t, H is the Heaviside function and a_f is the threshold of touch.

2) Accuracy: We calculate the accuracy of manipulation as the Euclidian distance of the actual object pose $\mathbf{x}_o$ to the desired object pose $\mathbf{x}_o^{des}$, averaged over total time T:

$$\sum_{t=0}^{T} \|\mathbf{x}_o^{des}(t) - \mathbf{x}_o(t)\|/T \tag{11}$$

We are reporting the position error e_{pos} in *cm* and the orientation error e_{rot} in *degrees*, separately. We calculate the orientation error as the angle between the unit x-axes of the actual and the desired rotation matrices of the object frame.

B. Experiment setup

We test the method on a real robot setup consisting of a 16-DoF Allegro robot hand mounted on a 6-DoF UR5 robot arm. The arm is used only for carrying the hand, it is not involved in the evaluated manipulation task. All 4 fingers of the robot are used for manipulation.

Fingertips of the robot hand are equipped with the Optoforce OMD 40 *N* force sensors. These sensors have 1000 *hz* frequency and a flexible structure. The sensors are used to detect contact pressures during grasping and manipulation. This information is used for safety during grasp and the calculation of the *contact stability* metric during manipulation.

All objects and the Allegro hand have ARUCO markers [69] on their front face (Fig. 6.b). These markers are tracked using a Kinect 2 camera that is facing the table where the experiments are carried out. We do not use any visual feedback in control, but we rely on ARUCO markers to acquire the object poses for the *accuracy* evaluation. The marker on the robot acts as a reference frame so that the object rotation is independent of the camera position. Our method is implemented as a set of *robot operating system (ROS)* packages. The codebase and the data of our experiments are available online[1]. The UR5 arm is controlled using *MoveIt!* library (*moveit.ros.org*). The Allegro hand inverse kinematics and manipulator Jacobian are calculated using *Orocos KDL* library (*www.orocos.org*). We use *DMPBBO* library for the implementation of the DMPs [70].

[1]https://gokhansolak.github.io/haptic-exploration-for-dexterous-manipulation

TABLE I
EXPERIMENT RESULTS: CONTACT LOSS AND TRACKING ERRORS ARE CALCULATED AS DESCRIBED IN SEC. V-A1 AND V-A2.

		contact loss (%)		orientation error e_{rot} (°)		position error e_{pos} (cm)		failures	
obj.	task	baseline	exploration	baseline	exploration	baseline	exploration	base.	expl.
1	1	3.752 ± 4.726	$\mathbf{0.000 \pm 0.000}$	6.126 ± 1.195	$\mathbf{5.301 \pm 0.505}$	0.561 ± 0.069	0.581 ± 0.087	0	0
	2	11.554 ± 8.469	$\mathbf{2.789 \pm 5.673}$	16.108 ± 2.249	$\mathbf{10.873 \pm 0.792}$	0.649 ± 0.074	0.645 ± 0.103	1	0
2	1	3.791 ± 9.135	0.006 ± 0.020	7.140 ± 3.623	5.558 ± 0.756	0.570 ± 0.075	0.590 ± 0.061	0	0
	2	13.935 ± 5.089	$\mathbf{2.725 \pm 4.462}$	17.494 ± 2.282	$\mathbf{10.263 \pm 0.778}$	0.760 ± 0.078	0.726 ± 0.087	2	0
3	1	1.176 ± 1.606	$\mathbf{0.000 \pm 0.000}$	6.839 ± 1.546	$\mathbf{4.950 \pm 0.884}$	0.598 ± 0.063	0.593 ± 0.099	0	0
	2	17.869 ± 9.100	$\mathbf{5.884 \pm 6.680}$	16.762 ± 1.797	$\mathbf{11.726 \pm 1.312}$	0.643 ± 0.095	0.660 ± 0.157	1	0

We used the academic version of *ALGLIB* (*www.alglib.net*) for optimisation. The non-linearly constrained optimisation problem (7) is solved using the *augmented Lagrangian* algorithm with analytical gradients. This algorithm may allow some violation of constraints. An instance of the problem is solved in 3-10 *ms*. The least-squares optimisation problem (6) is modelled as quadratic programming and solved using the *QuickQP* solver of *ALGLIB*.

C. Dataset

We evaluate our system by teaching two representative motions by demonstration (Fig. 6.b). Both are demonstrated kinesthetically, by holding the robot physically and guiding its fingers. Although both are unconstrained *6-DoF* trajectories, *task 1* is dominantly about translation and *task 2*, rotation. A video of the experiments is available[1].

The learned motions are then applied to objects with different shapes and friction properties. Our setup includes 3 objects of different shapes and textures as illustrated in Fig. 6.a. The objects are handcrafted from extruded polystyrene, hence their surfaces are not perfectly smooth. *Objects 1* and *3* are covered with tape, and *object 2* is covered with paper to have different friction properties. Different object shapes are chosen to have different local surface normals. People use objects that show these characteristics on daily basis, e.g., a bottle can have a convex, straight or concave facet.

D. Procedure

Each experiment begins with the object placed manually on a marked area. The robot hand is brought closer to the area by the robot arm, the hand takes a programmed pre-grasp shape and puts its fingertips on the object. We rely on pre-programmed grasp planning because it is out of our scope. This behaviour could also be achieved by a vision-based grasp planning algorithm [71].

When all the fingers are in contact with the object, haptic exploration begins. Haptic exploration proceeds in two stages. *Stage 1* probes the local surface orientation by applying the centre-bound forces. If the force is not incidentally upright the fingers slip along the surface. Using this observation, the forces are updated to be perpendicular to the slip vectors and the grasp is repeated as in Fig. 4.b. *Stage 1* lasts about *12 s*.

After estimating the surface normal, the exploration proceeds with *stage 2*. In this stage, the grasping forces are rotated individually to explore the 3-dimensional contact force space. It searches for the boundary of the friction cone iteratively (Fig. 5). *Stage 2* lasts about *130 s* on average. When the friction cone is estimated, the object is regrasped and the GFO initiates with the extracted information. The GFO continues until the end of the experiment.

After the exploration, the object is lifted by the robot arm. The hand attempts the manipulation task that is encoded by the learned DMPs. The learned motion primitives are modulated with the initial object pose to create an object trajectory. This trajectory is then mapped to each finger joint and executed using PID control. Both tasks are executed in *16 s*.

The baseline method does not involve the haptic exploration procedure. In that case, the object is grasped using the VSF as described in Sec. IV-F. After the grasp, manipulation is carried the same way as above.

E. Results

We have carried out the experiments for 2 tasks (Fig. 6.b) and 3 objects (Fig. 6.a) with and without haptic exploration. Each case is repeated 10 times.[2] Table I presents the *contact stability* and *accuracy* results. It also includes the failure counts for all cases where the object was lost during manipulation. The failed cases are discarded during the calculation of the performance metrics. For an easier comparison of the results, we also provide the bar plots in Fig. 7. The orientation trajectories are plotted for *task 2* executed on *object 1* to demonstrate the tracking accuracy over time (Fig. 8).

We mark the cases with significantly better outcomes using boldface fonts in the table. We used Welch's *t*-test for the significance analysis ($p < .05$) for its robustness when the compared populations have different variances [72].

F. Discussion

The results show that the exploration procedure improves the contact stability and the orientation accuracy significantly in most of the cases while the position accuracy does not show a significant difference overall (Fig. 7). The exploration also prevented failures that occurred in *task 2* with the baseline.

The inferior results of the baseline method can be attributed mainly to contact slips. As it is illustrated in Fig. 1, the VSF approach works best when the object has a round convex shape where the forces are incidentally applied upright to the surface of the object. In an unstructured environment, it is likely to encounter objects for which this is not the case.

[2]For the *object 1-task 1* case we considered only 9 repetitions, as one trial was discarded during data analysis due to faulty marker tracking.

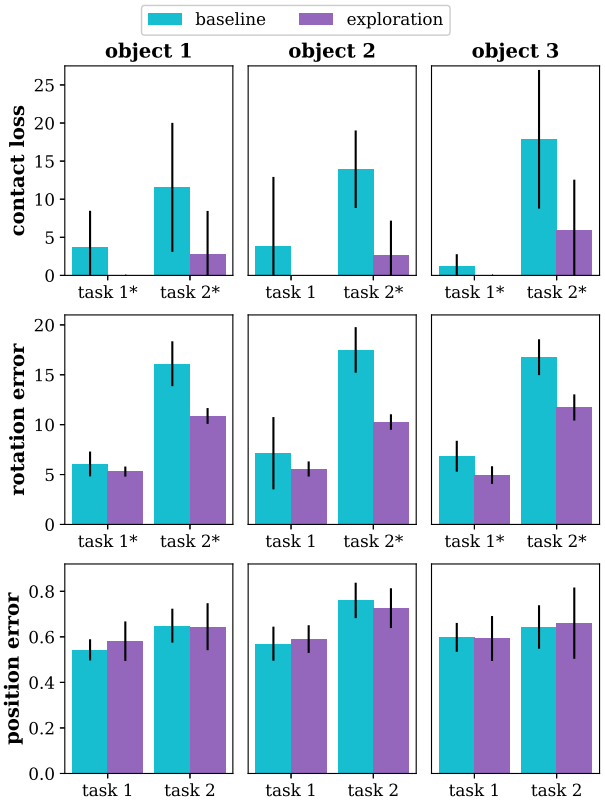

Fig. 7. Bar plots showing the performance of two approaches for each object-task case. Bars represent the mean of the given metric, lines represent the standard deviation. The rows present different metrics: *contact loss*, e_{rot}, e_{pos}, respectively. Significantly different ($p < .05$) results are marked with an asterisk.

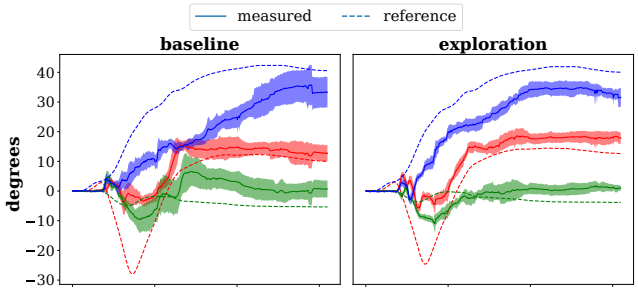

Fig. 8. Object trajectory tracking plots of *task 2* executed on *object 1* for orientations around axes x (red), y (green), and z (blue). The reference trajectories are shown with discrete lines. The actual trajectories are shown with solid lines, as the mean and standard deviation of the 10 repetitions. The standard deviations of the reference trajectories are not shown as they are very narrow.

The higher contact loss in the baseline outcomes can be explained by the observation that all contacts slide towards the centre, pushing and blocking each other. In some executions of *task 2*, the index and the ring fingers squeezed the middle finger away from the object during rotation. Additionally, we observe short temporal contact losses during the manipulation that can be attributed to the fact that the baseline applies more tangential forces than the proposed method. In some cases, we observed a partial loss of the grasp during manipulation where a fingertip slips out of the surface while the object is held with the other fingers. The exploration improves the contact stability because it adapts the contact forces to increase the normal components and satisfy the estimated friction cones.

The improvement in the orientation error e_{rot} can be explained by the observation that in the baseline all contacts slide towards the centre, hence it is hard to apply the desired torque on the object frame. On the other hand, the proposed method helps to maintain the initial contact points which are further from the centre, hence producing higher object torques with lower contact forces. Fig. 8 shows this improvement clearly: the exploration method both reduces the tracking error and increases the repeatability of the motion by reducing the deviation. The consistency of the performance can be attributed to the preservation of the intended grasp configuration. Please note that Fig. 8 shows the motion of the object frame as tracked using ARUCO markers; thus, the exact tracking error seen in this figure was not available to the controller.

The improvement of contact losses and rotation errors was minor for *task 1*. This is understandable as this task stays in safer parts of the workspace in comparison to *task 2*, and the rotational component of this motion is smaller.

We see very similar position errors e_{pos} overall. In all cases, the average error is around 6.5 *mm*. This is satisfactory, given the accuracy of the robot hand and the ARUCO marker tracking. The mean difference between the two approaches is less than 0.5 *mm*. The translation accuracy e_{pos} is not significantly affected by the contact slips because the translation direction is different from the possible finger slip directions.

VI. Conclusion

We presented a novel haptic exploration method to extract local surface normal and friction coefficient information at contact points to improve in-hand manipulation performance. Our haptic exploration method does not require complex special-purpose sensors and object shape information. It is designed to conclude in a short time without lifting the object so it can be applied safely in unstructured environments.

The haptic exploration method and our earlier object-centric LfD method constitutes a framework for in-hand manipulation of unknown objects. This framework answers the problems of the missing object and task information, towards the goal of in-hand manipulation in unstructured environments. It learns novel action knowledge from human demonstration and acquires object information by autonomous haptic exploration. We validate the contribution of haptic exploration directly on robotic in-hand manipulation performance which has not been shown before, to our knowledge.

In our previous work, we have used the VSF for in-hand manipulation of unknown objects. However, for some object shapes, grasp stability and manipulation accuracy issues arise. Haptic exploration is introduced to answer these issues. Grasping with online force optimisation that is informed by haptic exploration improves stability significantly. It also improves the manipulation accuracy, specifically the orientation error.

Our method is limited to rigid, high-DoF anthropomorphic robot hands. In the case of a non-rigid robot hand, our method would require further sensing capabilities to accurately measure the fingertip positions and feedforward force. High-DoF is required to apply the diverse grasping forces at different

contact points. Also, our haptic exploration approach is limited to rigid objects as the presumed slip behaviour does not occur for deformable objects.

VII. Acknowledgments

Work partially supported by the EPSRC UK through projects NCNR (EP/R02572X/1) and MAN^3 (EP/S00453X/1).

References

[1] C. C. Kemp, A. Edsinger, and E. Torres-Jara, "Challenges for robot manipulation in human environments [grand challenges of robotics]," *IEEE Robotics Automation Magazine*, vol. 14, no. 1, pp. 20–29, 2007.

[2] I. Vitanov, I. Farkhatdinov, B. Denoun, F. Palermo, A. Otaran, J. Brown, B. Omarali, T. Abrar, M. Hansard, C. Oh *et al.*, "A suite of robotic solutions for nuclear waste decommissioning," *Robotics*, vol. 10, no. 4, p. 112, 2021.

[3] D. Katz, J. Kenney, and O. Brock, "How can robots succeed in unstructured environments," in *In Workshop on Robot Manipulation: Intelligence in Human Environments at Robotics: Science and Systems*. Citeseer, 2008.

[4] I. M. Bullock, R. R. Ma, and A. M. Dollar, "A hand-centric classification of human and robot dexterous manipulation," *IEEE transactions on Haptics*, vol. 6, no. 2, pp. 129–144, 2012.

[5] G. Solak and L. Jamone, "Learning by demonstration and robust control of dexterous in-hand robotic manipulation skills," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2019.

[6] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: learning attractor models for motor behaviors," *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.

[7] T. Wimböck, C. Ott, A. Albu-Schäffer, and G. Hirzinger, "Comparison of object-level grasp controllers for dynamic dexterous manipulation," *The International Journal of Robotics Research*, vol. 31, no. 1, pp. 3–23, 2012.

[8] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, vol. 48, pp. 54–67, 2017.

[9] Q. Li, O. Kroemer, Z. Su, F. F. Veiga, M. Kaboli, and H. J. Ritter, "A review of tactile information: Perception and action through touch," *IEEE Transactions on Robotics*, vol. 36, no. 6, pp. 1619–1634, 2020.

[10] B. Sundaralingam and T. Hermans, "In-hand object-dynamics inference using tactile fingertips," *IEEE Transactions on Robotics*, 2021.

[11] J. A. Fishel and G. E. Loeb, "Bayesian exploration for intelligent identification of textures," *Frontiers in neurorobotics*, vol. 6, p. 4, 2012.

[12] N. F. Lepora, U. Martinez-Hernandez, and T. J. Prescott, "Active touch for robust perception under position uncertainty," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 3020–3025.

[13] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, "Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 297–303.

[14] Z. Deng, Y. Jonetzko, L. Zhang, and J. Zhang, "Grasping force control of multi-fingered robotic hands through tactile sensing for object stabilization," *Sensors*, vol. 20, no. 4, p. 1050, 2020.

[15] K. Tahara, S. Arimoto, and M. Yoshida, "Dynamic object manipulation using a virtual frame by a triple soft-fingered robotic hand," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 4322–4327.

[16] M. Li, H. Yin, K. Tahara, and A. Billard, "Learning object-level impedance control for robust grasping and dexterous manipulation," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 6784–6791.

[17] M. Li, Y. Bekiroglu, D. Kragic, and A. Billard, "Learning of grasp adaptation through experience and tactile sensing," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. Ieee, 2014, pp. 3339–3346.

[18] Z. Chen, C. Ott, and N. Y. Lii, "A compliant multi-finger grasp approach control strategy based on the virtual spring framework," in *International Conference on Intelligent Robotics and Applications*. Springer, 2015, pp. 381–395.

[19] G. Solak and L. Jamone, "Haptic exploration for dexterous manipulation stability experiments," Aug 2021. [Online]. Available: https://doi.org/10.6084/m9.figshare.c.5560167.v2

[20] F. Veiga, J. Peters, and T. Hermans, "Grip stabilization of novel objects using slip prediction," *IEEE Transactions on Haptics*, vol. 11, no. 4, pp. 531–542, 2018.

[21] N. Sommer and A. Billard, "Multi-contact haptic exploration and grasping with tactile sensors," *Robotics and autonomous systems*, vol. 85, pp. 48–61, 2016.

[22] M. Kaboli, K. Yao, D. Feng, and G. Cheng, "Tactile-based active object discrimination and target object search in an unknown workspace," *Autonomous Robots*, vol. 43, no. 1, pp. 123–152, 2019.

[23] M. S. Siddiqui, C. Coppola, G. Solak, and L. Jamone, "Grasp stability prediction for a dexterous robotic hand combining depth vision and haptic bayesian exploration," *Frontiers in Robotics and AI*, vol. 8, p. 237, 2021.

[24] C. Pan, M. Lepert, S. Yuan, R. Antonova, and J. Bohg, "Task-driven in-hand manipulation of unknown objects with tactile sensing," *arXiv preprint arXiv:2210.13403*, 2022.

[25] G. Averta, F. Barontini, I. Valdambrini, P. Cheli, D. Bacciu, and M. Bianchi, "Learning to prevent grasp failure with soft hands: From online prediction to dual-arm grasp recovery," *Advanced Intelligent Systems*, vol. 4, no. 3, p. 2100146, 2022.

[26] L. Deng, Y. Shen, G. Fan, X. He, Z. Li, and Y. Yuan, "Design of a soft gripper with improved microfluidic tactile sensors for classification of deformable objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5607–5614, 2022.

[27] W. Bottcher, P. Machado, N. Lama, and T. M. McGinnity, "Object recognition for robotics from tactile time series data utilising different neural network architectures," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–8.

[28] B. Wu, I. Akinola, J. Varley, and P. K. Allen, "Mat: Multi-fingered adaptive tactile grasping via deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2020, pp. 142–161.

[29] C. de Farias, N. Marturi, R. Stolkin, and Y. Bekiroglu, "Simultaneous tactile exploration and grasp refinement for unknown objects," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3349–3356, 2021.

[30] H. Dang and P. K. Allen, "Grasp adjustment on novel objects using tactile experience from similar local geometry," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4007–4012.

[31] N. Wettels and G. E. Loeb, "Haptic feature extraction from a biomimetic tactile sensor: force, contact location and curvature," in *2011 IEEE International Conference on Robotics and Biomimetics*. IEEE, 2011, pp. 2471–2478.

[32] W. Chen, H. Khamis, I. Birznieks, N. F. Lepora, and S. J. Redmond, "Tactile sensors for friction estimation and incipient slip detection—toward dexterous robotic manipulation: A review," *IEEE Sensors Journal*, vol. 18, no. 22, pp. 9049–9064, 2018.

[33] R. A. Romeo and L. Zollo, "Methods and sensors for slip detection in robotics: A survey," *IEEE Access*, vol. 8, pp. 73 027–73 050, 2020.

[34] T. Okatani, H. Takahashi, K. Noda, T. Takahata, K. Matsumoto, and I. Shimoyama, "A tactile sensor using piezoresistive beams for detection of the coefficient of static friction," *Sensors*, vol. 16, no. 5, p. 718, 2016.

[35] W. Chen, S. Rodpongpun, W. Luo, N. Isaacson, L. Kark, H. Khamis, and S. J. Redmond, "An eight-legged tactile sensor to estimate coefficient of static friction," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2015, pp. 4407–4410.

[36] H. Khamis, R. I. Albero, M. Salerno, A. S. Idil, A. Loizou, and S. J. Redmond, "Papillarray: An incipient slip sensor for dexterous robotic or prosthetic manipulation–design and prototype validation," *Sensors and Actuators A: Physical*, vol. 270, pp. 195–204, 2018.

[37] K. Van Wyk and J. Falco, "Calibration and analysis of tactile sensors as slip detectors," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2744–2751.

[38] R. Fernandez, I. Payo, A. Vazquez, and J. Becedas, "Micro-vibration-based slip detection in tactile force sensors," *Sensors*, vol. 14, no. 1, pp. 709–730, 2014.

[39] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson, "Measurement of shear and slip with a gelsight tactile sensor," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 304–311.

[40] G. De Maria, P. Falco, C. Natale, and S. Pirozzi, "Integrated force/tactile sensing: The enabling technology for slipping detection and avoidance," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3883–3889.

[41] R. Zenha, B. Denoun, C. Coppola, and L. Jamone, "Tactile slip detection in the wild leveraging distributed sensing of both normal and shear forces," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 2708–2713.

[42] J. Reinecke, A. Dietrich, F. Schmidt, and M. Chalon, "Experimental comparison of slip detection strategies by tactile sensing with the biotac® on the dlr hand arm system," in *2014 IEEE international Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2742–2748.

[43] L. Jamone, L. Natale, G. Metta, and G. Sandini, "Highly sensitive soft tactile sensors for an anthropomorphic robotic hand," *IEEE Sensors Journal*, vol. 15, no. 8, pp. 4226–4233, 2015.

[44] T. Paulino, P. Ribeiro, M. Neto, S. Cardoso, A. Schmitz, J. Santos-Victor, A. Bernardino, and L. Jamone, "Low-cost 3-axis soft tactile sensors for the human-friendly robot vizzy," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 966–971.

[45] T. P. Tomo, A. Schmitz, W. K. Wong, H. Kristanto, S. Somlor, J. Hwang, L. Jamone, and S. Sugano, "Covering a robot fingertip with uskin: A soft electronic skin with distributed 3-axis force sensitive elements for robot hands," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 124–131, 2018.

[46] T. P. Tomo, M. Regoli, A. Schmitz, L. Natale, H. Kristanto, S. Somlor, L. Jamone, G. Metta, and S. Sugano, "A new silicone structure for uskin—a soft, distributed, digital 3-axis skin sensor and its integration on the humanoid robot icub," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2584–2591, 2018.

[47] D. Prattichizzo, M. Pozzi, and M. Malvezzi, "Dexterous manipulation," *Encyclopedia of Robotics. Springer Berlin Heidelberg, Berlin, Heidelberg*, pp. 1–8, 2020.

[48] A. M. Okamura, N. Smaby, and M. R. Cutkosky, "An overview of dexterous manipulation," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 255–262.

[49] N. Chavan-Dafle, R. Holladay, and A. Rodriguez, "Planar in-hand manipulation via motion cones," *International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 163–182, 2020.

[50] B. Calli, K. Srinivasan, A. Morgan, and A. M. Dollar, "Learning modes of within-hand manipulation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3145–3151.

[51] R. Ozawa and K. Tahara, "Grasp and dexterous manipulation of multi-fingered robotic hands: a review from a control view point," *Advanced Robotics*, vol. 31, no. 19-20, pp. 1030–1050, 2017.

[52] P. Michelman, "Precision object manipulation with a multifingered robot hand," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 1, pp. 105–113, 1998.

[53] Z. Kappassov, J.-A. Corrales, and V. Perdereau, "Tactile sensing in dexterous robot hands," *Robotics and Autonomous Systems*, vol. 74, pp. 195–220, 2015.

[54] M. Pfanne, M. Chalon, F. Stulp, H. Ritter, and A. Albu-Schäffer, "Object-level impedance control for dexterous in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2987–2994, 2020.

[55] A. Bhatt, A. Sieler, S. Puhlmann, and O. Brock, "Surprisingly robust in-hand manipulation: An empirical study," *arXiv preprint arXiv:2201.11503*, 2022.

[56] A. S. Morgan, K. Hang, B. Wen, K. Bekris, and A. M. Dollar, "Complex in-hand manipulation via compliance-enabled finger gaiting and multimodal planning," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4821–4828, 2022.

[57] T. Wimboeck, C. Ott, and G. Hirzinger, "Passivity-based object-level impedance control for a multifingered hand," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 4621–4627.

[58] V. Ruiz Garate, M. Pozzi, D. Prattichizzo, and A. Ajoudani, "A bio-inspired grasp stiffness control for robotic hands," *Frontiers in Robotics and AI*, vol. 5, p. 89, 2018.

[59] W. Shaw-Cortez, D. Oetomo, C. Manzie, and P. Choong, "Tactile-based blind grasping: A discrete-time object manipulation controller for robotic hands," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1064–1071, 2018.

[60] B. Sundaralingam and T. Hermans, "Relaxed-rigidity constraints: kinematic trajectory optimization and collision avoidance for in-grasp manipulation," *Autonomous Robots*, vol. 43, no. 2, pp. 469–483, 2019.

[61] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," in *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018.

[62] M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.

[63] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112.

[64] A. Gupta, J. Yu, T. Z. Zhao, V. Kumar, A. Rovinsky, K. Xu, T. Devlin, and S. Levine, "Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 6664–6671.

[65] A. Cloutier and J. Yang, "Grasping force optimization approaches for anthropomorphic hands," *Journal of Mechanisms and Robotics*, vol. 10, no. 1, 2018.

[66] F.-T. Cheng and D. E. Orin, "Efficient algorithm for optimal force distribution-the compact-dual lp method," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 2, pp. 178–187, 1990.

[67] S. P. Boyd and B. Wegbreit, "Fast computation of optimal contact forces," *IEEE Transactions on Robotics*, vol. 23, no. 6, pp. 1117–1132, 2007.

[68] A. Ude, B. Nemec, T. Petrić, and J. Morimoto, "Orientation in cartesian space dynamic movement primitives," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 2997–3004.

[69] S. Garrido-Jurado, R. Muñoz-Salinas, F. Madrid-Cuevas, and M. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280 – 2292, 2014.

[70] F. Stulp, "Dmpbbo–a c++ library for black-box optimization of dynamical movement primitives," 2014.

[71] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2013.

[72] M. Delacre, D. Lakens, and C. Leys, "Why psychologists should by default use welch's t-test instead of student's t-test," *International Review of Social Psychology*, vol. 30, no. 1, 2017.

Gokhan Solak received a B.S. degree in electrical and electronics engineering from Anadolu University, Eskisehir, Turkey, in 2013. He obtained an M.Sc. degree in computer engineering from Istanbul Technical University, Istanbul, Turkey, in 2017. He obtained his Ph.D. degree from computer science at Queen Mary University of London, UK, in 2022. Currently, he is working as a postdoctoral researcher at the Italian Institute of Technology, Genoa. His research interests include robotic grasping and manipulation, robot learning, and tactile sensing.

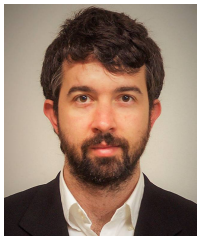

Lorenzo Jamone (S'06-M'11) is a Senior Lecturer in Robotics at the Queen Mary University of London (UK) where he is the founder and director of the CRISP group: Cognitive Robotics and Intelligent Systems for the People. He received the MS degree (honours) in computer engineering from the University of Genoa, Genoa, Italy, in 2006, and the PhD degree in humanoid technologies from the University of Genoa (Italy) and the Italian Institute of Technology, in 2010. He was an Associate Researcher at the Takanishi Laboratory, Waseda University (Tokyo, Japan) from 2010 to 2012, and at the Computer and Robot Vision Laboratory, Instituto Superior Técnico (Lisbon, Portugal) from 2012 to 2016. He has over 90 publications with an H-index of 26. His current research interests include cognitive robotics, robotic manipulation, force and tactile sensing.