

Inferring the Effects of Wiping Motions based on Haptic Perception

Daniel Leidner¹ and Michael Beetz²

Abstract—Future service robots are expected to achieve high-quality task performance for everyday household chores. Some of the most frequent tasks in this domain are related to wiping of surfaces, such as vacuuming the floor, sweeping dust, or cleaning windows. However, the performance for these tasks is not directly observable as small dirt particles, dust, and residual water are hardly perceivable by means of computer vision. In this work we propose to utilize haptic perception paired with a qualitative effect representation to reason about the task performance of robotic wiping motions despite poor visual information. In particular, we relate the desired contact force to the measured end-effector force in order to simulate the effect of previously executed wiping motions. This way we are not just able to distinguish good from bad contact situations, but also replan recovery motions w.r.t. the effect-space to accomplish the commanded cleaning task subsequently. We evaluate our approach in a set of experiments with the robot Rollin’ Justin.

I. INTRODUCTION

Service robots are envisaged as universal assistant in everyday environments. As such, they will have to cope with a wide variety of daily household chores, including cooking, organization and cleaning. Within the context of robotic manipulation, but especially for cleaning tasks, it is important to monitor the performance of the executed actions. Typically, humans evaluate the outcome of their actions based on visual perception. For example, collecting shattered shards of a broken mug with a broom will eventually result in accumulated pile of shards visually noticeable at the dedicated goal region. Another example is the absorption of larger dirt particles with a vacuum cleaner, which results in a clean surface if all particles are removed. Based on these assumptions, there have been some efforts on effect oriented robotic cleaning in the past [1], [2], [3]. However, visual feedback is an unreliable source of information to evaluate some of the most frequent cleaning tasks including vacuuming the floor, dusting surfaces, and window wiping, as small dirt particles, dust, and streaks of water are hardly perceivable in camera images or depth images, especially on transparent and reflecting surfaces (see Fig. 1).

To overcome this issue, we argue that according to research on neurobiology [4] humans do not solely rely on perceptual feedback, but also maintain knowledge of their manipulation actions and the resulting effects in form of abstract process models. For example, in order to clean a sideboard with a feather duster, a human would try to cover the entire surface of the target object by wiping along it with



Fig. 1. Reflections and hardly perceivable streaks of detergent impede visual perception for the tasks like cleaning windows or solar panels.

the tool. This behavior is based on the knowledge that dust particles are electrostatically absorbed by the feather duster and the assumption that the dust is equally distributed on the planar surface. Consequently, humans are able to infer that the desired effect (i.e. having the side board cleaned from dust) is successfully accomplished after the tool has been in physical contact with the whole target area.

Nevertheless, the removal of the dust is not visible and the effect is not directly perceivable. The only reliable feedback left is the haptic information occurring from the contact during the wiping motions. The sense of touch is an essential factor for the task reasoning and effect inference of humans in the absence of vision [5]. Haptic feedback provides humans not only with the information that contact was successfully established, but also provides the basis to rate the task performance and even detect performance errors based on the comparison of the desired contact forces and the actual sensed force [6]. In case of bad contact introduced by stick-slip or uneven areas (recognizable as lower resistance force) humans may furthermore decide to revisit the affected regions to improve the cleaning result.

To this end, the combination of cognitive capabilities and haptic perception enables humans to qualitatively reason about the effects of their motions and solve even complex cleaning tasks despite poor visual perception. Accordingly, we propose to utilize the torque sensing capabilities of compliant light weight robots [7] to infer contact situations of compliant wiping motions and measure the task performance based on a qualitative effect model.

In particular, this work is based on the representation and planning methods for wiping tasks introduced in [8]. The contributions in this work include (i) an approach to detect contact during real world wiping motions, (ii) inference methods to estimate the performance of these motions,

¹ Institute of Robotics and Mechatronics, German Aerospace Center (DLR), Wessling, Germany. Contact: daniel.leidner@dlr.de

² Institute for Artificial Intelligence, University of Bremen, Bremen, Germany.

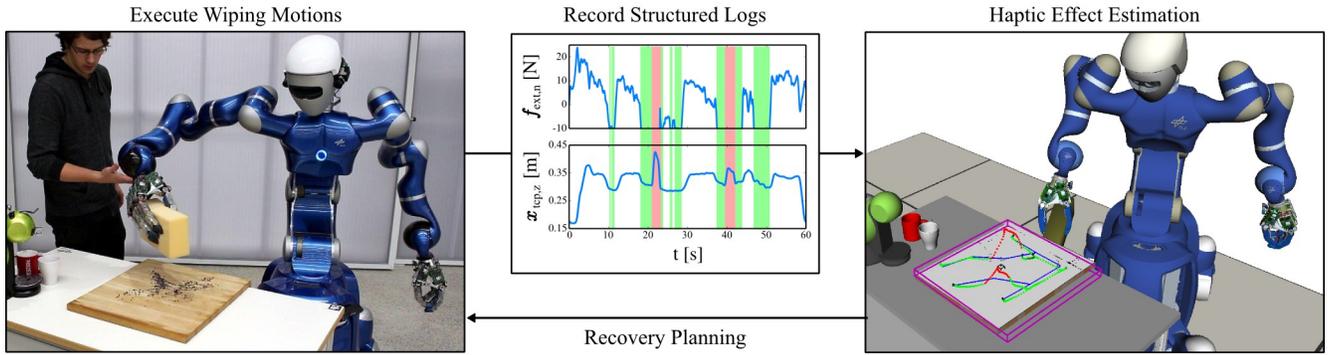


Fig. 2. Rollin’ Justin is disturbed while executing wiping motions (left). End-effector positions and forces (top) are logged to rate the performance based on a particle simulation (right). The outcome of this qualitative estimation is used to plan recovery motions to enhance the performance afterwards (bottom).

(iii) a probabilistic contact model to incorporate different contact situations arising from different tool-medium-surface constellations, and (iv) an approach to distinguish “good” contact situations from “bad” contact situations. By utilizing the planning algorithms introduced in [8], (v) we are further able to replan additional wiping motions to enhance the cleaning result w.r.t. bad contact situations introduced by external disturbances. The effect inference is realized based on real world measurements recorded during experiments conducted with the humanoid robot Rollin’ Justin [9].

The remainder of this work is structured as follows. After a review of the state-of-the-art, we introduce our particle distribution model and outline the planning algorithms to compute efficient cleaning motions in Sec. III. Based on this representation, we describe our effect inference strategy to estimate the task performance of wiping motions in a cleaning scenario in Sec. IV. Eventually, we present our approach on failure detection and recovery in Sec. V which serves also as evaluation of our proposed methods.

II. RELATED WORK

Wiping motions are often considered as fundamental part of cleaning actions and have therefore been investigated in robotics research in the recent past to some extent. The articles most related to our work are listed as follows. Gams *et al.* [10] investigate cleaning motions from a learning point of view. They exploit the compliance of a LWR III robot to adapt to unknown surface geometries and modify the cleaning motion w.r.t. physical contact introduced by a human tutor. This way a human can directly modify the periodic cleaning motions. However, there is no particular goal specified as the wiping motions are considered as prototypical actions. Hess *et al.* conducted research on robotic cleaning in a series of papers [1], [11]. They investigate cleaning as a path coverage problem for robotic manipulators and vacuum robots. In [1] they learn the effect of a vacuum cleaner moving along a planar surface by utilizing visual feedback based on color segmentation. The robot can enhance the task execution for future trials as it generates motions that cover only the dirty areas. A similar approach is used to plan optimal motions for a vacuuming robot [11]. Based on a discretized dirt distribution grid model the robot plans the motions to clean the floor most efficiently. This model is

learned by utilizing a dirt sensor that measures the impact of dirt particles. In both articles it is implicitly assumed that dirt is absorbed upon contact. Martinez *et al.* [2] investigate planning for robotic cleaning by wiping with a sponge under the assumption that the particles are pushed upon contact. Additionally, the authors propose to represent dirt accumulations as ellipses to have them accessible as semantic predicates [12] for automated planning [13]. The ellipses are computed based on color image data recorded after each wiping motion. Similarly, Do *et al.* [3] utilize a perceptual representation of dirt distributions on a target surface to derive a scalar value to rate the task performance w.r.t. different object properties and action parameters. However, even though these research groups make implicit assumptions of the wiping effect, there is no underlying model available that could be used to predict the actual task performance without visual validation as physical contact is not explicitly modeled. Accordingly, the robots cannot make assumptions on the task performance from haptic feedback.

In contrast, there has been some work on contact effect modeling for wiping motions and effect estimation for robotic manipulation tasks in general. Kunze *et al.* [14] utilized simplified process model based on particle simulations to infer the effect of tools interacting with their environment. In particular, they simulate the effect of a sponge absorbing liquids in contact. This way, a qualitative effect inference can be conducted based on the absorbed and leftover water particles. Winkler *et al.* [15] maintain expectation about the outcome of planned manipulations in pick-and-place scenarios. Based on observation of relevant task parameters (e.g. gripper forces during object transitions) a robot can learn when an action was successfully executed or failed. Pastor *et al.* [16] propose to learn motor skills in form of Dynamic Movement Primitives. Additionally they predict the task outcome of the manipulation tasks based on statistical hypothesis testing w.r.t. low-level sensor streams which include force information recorded with fingertip pressure sensors. The proposed approach enables the robot to predict failure situations online. Obviously, contact force is also of interest for the contact-rich task of cleaning, however, to the authors knowledge, there has not yet been any efforts to qualitatively model the effect of wiping motions on a high level of abstraction based on low-level control feedback.

This paper builds on the findings of our earlier work, where we developed an approach to reason about wiping motions and their effects based on a particle distribution model [8]. This model was initially developed to plan tool motions in the effect-space according to a given semantic goal (e. g. absorb the particles, collect the particles, or skim the particles away). By simulating the effect of the particle distribution w. r. t. the planned tool motions and the resulting contact we were able to *predict* the task outcome for the later execution of the motions. A brief experiment was conducted to evaluate the planning method.

We observed that the simulated expectation matched the real task outcome which motivates the work presented here. In particular, we aim to utilize the introduced particle representation to *estimate* the task performance of real world wiping motions. As a result we are eventually able to close the loop from semantic reasoning to low level control and vice versa as outlined in Fig. 2. In a nutshell, the robot executes wiping motions according to the planning methods introduced in [8] (left). During execution, we record the Cartesian end-effector position and the end-effector force (top, center). The measurements that match the desired force profile are considered in contact with the target surface (green bars). Whenever contact is detected, that is not aligned with the surface, it is considered as disturbance (red bars). Based on this information we simulate the effect to the particle distribution in order to evaluate the real world task outcome (right). The resulting particle distribution can be used to reinitialize the planning methods and compensate for execution errors (bottom, center).

III. EFFECT REPRESENTATION

As our long term goal is to develop a generalized framework for compliant robotic manipulation we have developed a classification of compliant manipulation tasks along with an abstract classification of wiping tasks in one of our previous works [17]. In particular, we proposed to classify compliant manipulation tasks according to the semantic contact situation between objects and the environment. Wiping actions are thereby considered as the interaction between tools, a surface, and most importantly the medium to be manipulated in form of particles or liquids, e. g. dirt, water, shards, paint, or other materials of daily living or manufacturing. Accordingly, we developed a particle representation to model the effect on the medium to autonomously plan wiping motions based on the desired semantic goal [8]. As we utilize the same model in this work to estimate the real world task outcome of wiping motions, we will briefly outline the representation and the motion planning algorithms respectively. For a comprehensive description of these methods please refer to [8].

A particle distribution model (Fig. 3, top left) is the basis for our motion planning algorithm that embodies a region coverage problem. Considering the distribution of particles

$$P = \left\{ (x_1, y_1), (x_2, y_2), \dots, (x_N, y_N) \mid x_i, y_i \in \mathbb{R} \wedge x_{\min} \leq x_i \leq x_{\max} \wedge y_{\min} \leq y_i \leq y_{\max} \right\}, \quad (1)$$

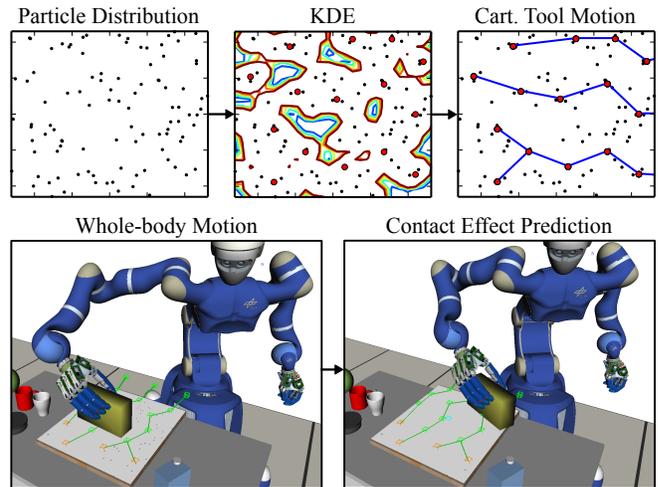


Fig. 3. Particle representation and reasoning methods developed in [8]. The waypoints (red dots) for the Cartesian tool motion are distributed based on high particle density areas according to the KDE. Based on these waypoints the task motion is computed (blue line) w. r. t. the semantic goal and grounded to whole-body joint motions. The simulated contact is used to estimate the wiping effect, recognizable in the absence of particles in the lower right sub-figure.

we proposed a graph-based reasoning algorithm to plan Cartesian wiping motions w. r. t. a desired semantic goal. The nodes of these graphs represent waypoints of Cartesian wiping motions. These nodes are distributed w. r. t. the extension of the particle model which can be detected by a vision algorithm, or estimated as a unified distribution on the target surface if the medium is hardly visible (e. g. dust). We have investigated several different methods to construct the graph structure, including a grid-based approach, Rapidly Exploring Random Trees (RRT), and Kernel Density Estimation (KDE) (Fig. 3, top center). The initial node distribution is the basis to grow the edges of the graph according to the semantic goal, which is a skim action in this example (Fig. 3, top right). Eventually, an inverse-kinematics-based motion planner is used to generate whole-body joint motions.

The particle distribution is projected on the planar target surface to be cleaned by the robot. By simulating the contact behavior between a tool and the particles we are able to estimate the effect of different wiping actions as illustrated in the bottom of Fig. 3. The contact model considers the volumetric model of the tool and the properties and position of each particle. For example, if a sponge is simulated to touch liquids, the resulting effect is the absorption of the liquid, i. e. the deletion of the particles. In case of a solid medium we simulate the accumulation of the particles as it is observed for collecting breadcrumbs at a dedicated destination or skimming breadcrumbs off the target surface. Accordingly, the particle distribution can be utilized to simulate the desired semantic *state change* $S_{t_0} \rightarrow S_g$ represented as PDDL predicate [12] (e. g. (*absorbed* ?m - medium ?s - surface), (*collected* ?m - medium ?s - surface), (*skimmed* ?m - medium ?s - surface)).

As visualized in Fig. 3, we have for now used the particle distribution to predict the wiping effect and the task

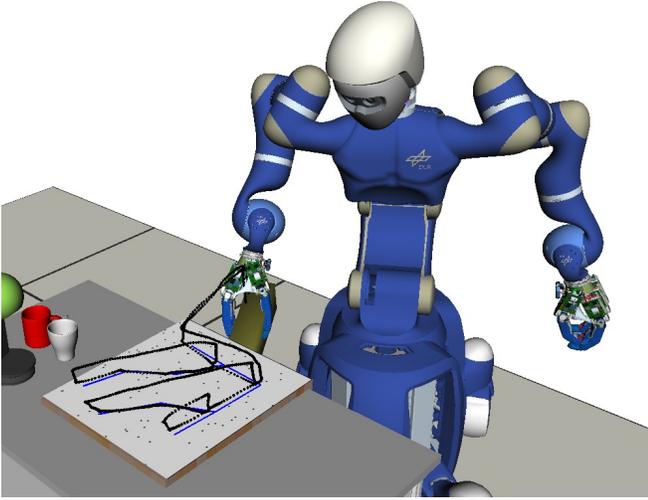


Fig. 4. The recorded Cartesian wiping motion of the robotic manipulator, i. e. the TCP of the sponge it is holding respectively (black dotted line) and the desired wiping motion in contact (blue lines).

performance based on simulated motions. A brief experiment with the real robot showed that this prediction matches the real world outcome which motivated the approach presented here. In comparison to [8], this work aims to estimate the task performance of real world wiping motions based on force information rather than predicting the effect beforehand based on purely kinematic simulations. To do so, we record telemetry data of the robots motions during wiping actions. This includes end-effector forces calculated from joint torque sensor measurements, as well as the measured Cartesian end-effector position. With this information we are able to reproduce the executed motions in simulation to estimate the contact between the tool and the particle distribution w.r.t. real world sensor readings and compute the real task performance respectively. The inference methods are described in the following sections.

IV. EFFECT INFERENCE

After the wiping motions are planned with the introduced algorithms, they are executed by means of a compliant whole-body impedance control framework [18], [19]. These motions are parameterized according to the task requirements which are provided by an object centric knowledge-base. Among others, the applied Cartesian force, the Cartesian stiffness, and the task hierarchy for the whole-body impedance are parameterized according to the requirements of the tool, the surface and the manipulated medium. Please note that we do not apply a force-control strategy, but rather exploit the compliant behavior to command a desired position *into* the target surface and saturate the maximum resistance force to the desired contact force [19]. As a consequence our control strategy allows for safe physical human-robot interaction during task execution. However, as the controller is not designed to adapt the Cartesian position w. r. t. the force measurements (unlike a hybrid position-force controller as described in [20]), it may happen that the robot

does not touch the surface appropriately in case of external disturbances. A possible solution to this will be discussed in Sec. V. The nominal case is discussed in this section.

To infer the effect of the robot motions, we record the actual measured end-effector positions \mathbf{x}_{act} and the controller forces at the end-effector \mathbf{f}_c for each timestamp i at 1 kHz. The controller force \mathbf{f}_c resulting from the commanded torque is thereby saturated to satisfy the desired contact

$$\boldsymbol{\tau} = -\mathbf{J}^T \underbrace{\mathcal{S} \left(\left(\frac{\partial V}{\partial \mathbf{x}} \right)^T + \mathbf{D}\dot{\mathbf{x}} \right)}_{\mathbf{f}_c}, \quad (2)$$

where \mathcal{S} describes the saturation function for the force terms $\mathbf{f}_{\text{spring}} = (\partial V / \partial \mathbf{x})^T$ plus $\mathbf{f}_{\text{damper}} = \mathbf{D}\dot{\mathbf{x}}$. Assuming a static case with negligible model uncertainty, the external contact force \mathbf{f}_{ext} counteracts the controller force \mathbf{f}_c . A sample trajectory of the *Tool Center Point* (TCP) is visualized as black dotted line in Fig. 4. The transformation of the TCP is thereby defined as $\mathbf{H}_{\text{tcp},i} = \mathbf{H}_{\text{act},i} \cdot \mathbf{H}_{\text{grasp}}^{-1}$, where the homogeneous transformation matrices \mathbf{H}_* correspond to the respective task space coordinates \mathbf{x}_* . We assume that the grasp transformation matrix $\mathbf{H}_{\text{grasp}}$ is constant during the task execution which is a valid assumption as the tools align to the curvature of the robot hand when utilizing firm power grasps. The blue lines on the target surface indicate the desired wiping motions of the sponge TCP in contact with the chopping board. The visualized motion is the outcome of a planned collect action, coincidentally forming an arrow-shaped path pointing towards the goal position for the particles.

For each measurement, we relate the recorded Cartesian position of the TCP to a contact force as it is visualized in Fig. 2. This way we are able to infer the segments of motion that are most likely in contact with the target surface. In the example at hand we are only interested in the force normal to the target surface $\mathbf{f}_{\text{ext},n}$. In particular, we analyze the normalized contact force

$$\mathbf{f}'_{n,i} = \frac{\min(\mathbf{f}_{\text{ext},n,\text{max}}, \mathbf{f}_{\text{ext},n,i})}{\mathbf{f}_{\text{ext},n,\text{min}}}, \quad (3)$$

for each force sample $\mathbf{f}_{\text{ext},n,i}$.

To estimate the positions in contact with the target surface in the first place, we do only consider the measured positions that show a high confidence for contact with the chopping board, illustrated as green dots on the left of Fig. 5. That is, we select only the positions

$$\mathbf{x}'_{\text{tcp},i} = \{ \mathbf{x}_{\text{tcp},i} \mid 0.9 \leq \mathbf{f}'_{n,i} \leq 1.0 \}, \quad (4)$$

which show a normalized force value of [0.9, 1.0], i. e. the measurements that show only 10% deviation from the desired contact force. As illustrated, most of the measurements close to the surface match these constraint and resemble the desired wiping trajectory. Even though this is already a quite accurate estimate of the contact motion, some segments were omitted due to lower contact forces introduced by friction effects.

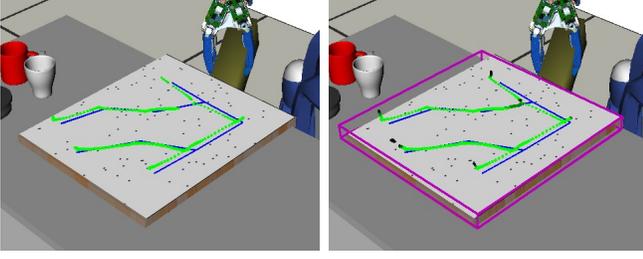


Fig. 5. Left: Closeup view of all measurements with high contact confidence $\mathbf{x}'_{\text{tcp},i}$ (green dots). These measurements are the basis for the target surface estimation. Right: The estimated target surface visualized as purple box. All positions of the sponge within this box are colored in green. The brighter the green color, the higher is the normalized contact force $\mathbf{f}'_{n,i}$.

To incorporate these left out yet still contact-rich segments, we propose to apply the *Random Sample Consensus (RanSaC)* algorithm [21] to estimate the transformation of the target surface. The RanSaC algorithm is commonly used in research on computer vision for plane fitting of surface elements perceived in visual data, e. g. point clouds computed based on RGB-D images. In this work we utilize the Cartesian TCP positions with high contact confidence $\mathbf{x}'_{\text{tcp},i}$ as dataset for the RanSaC algorithm. This allows a robot to estimate the target surface even under poor lighting conditions as they occur on reflecting and transparent glass panes. In particular,

$$\mathbf{x}_{\text{plane},i} = \text{RanSaC} \left(\sum_{i=0}^N \mathbf{x}'_{\text{tcp},i}, \epsilon \right), \quad (5)$$

where $\mathbf{x}_{\text{plane},i}$ constitute the measurements inside the estimated surface plane visualized as purple box in the right of Fig. 5. They are visualized as green dots of different brightness. A bright green color represents a high normalized contact force $\mathbf{f}'_{n,i}$, where as darker green colors (eventually fading to black) are representative for lower normalized contact forces $\mathbf{f}'_{n,i}$. The measurement points $\mathbf{x}_{\text{plane},i}$ are now considered in contact with the surface and are utilized to estimate the effect as described as follows.

Since lower contact forces may indicate a bad contact situation, we propose to integrate the normalized contact force $\mathbf{f}'_{n,i}$ to model the effect of wiping motions. In our previous work [8] we have predicted the effect of wiping motions based on simulated data obtained by means of a kinematics simulation. The contact model was limited to the volumetric model of the tool in relation to the position of the particles distributed on the target surface introduced in Sec. III. There was no notion of forces arising from contact involved in the simulated effect model.

By re-executing the recorded motions in simulation we are able to infer the real world effect of the prior executed wiping motions under consideration of the recorded contact forces. Our assumption is that contact forces close to the desired force have higher probability to produce the desired wiping effect to the medium. However, the actual effect depends on the properties of the tool-medium-surface tuple as described

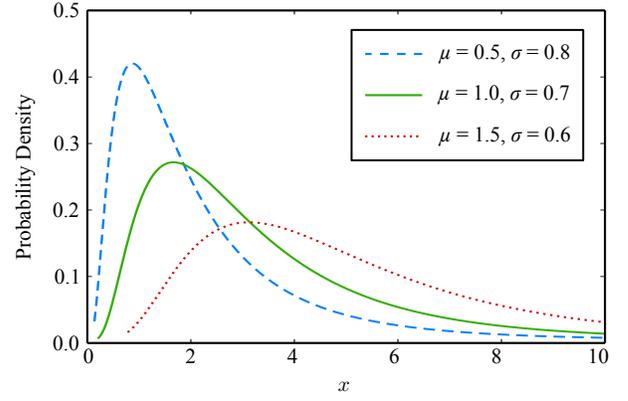


Fig. 6. Exemplary probability density functions for different tool-medium-surface constellations.

in our previous investigations on compliant manipulation tasks [17]. Among others, the exact tool geometry, the tool elasticity, the surface friction, the surface inclination, the medium friction, and the medium size and geometry influence the outcome of wiping actions. Most of these parameters are hardly assessable. Instead, we utilize the *log-likelihood function* $\log(\mathcal{L}(\theta|x))$ to simulate the effect on a qualitative basis.

- If $\mathbf{f}'_{n,i} \geq \|\log(\mathcal{L}(\theta|x_i))\|$, we simulate the contact behavior (i. e. move the affected particles parallel to the tool motion in case of collect and skim actions).
- If $\mathbf{f}'_{n,i} < \|\log(\mathcal{L}(\theta|x_i))\|$, we skip the simulation step and immediately proceed with the next measurement.

The logarithmic *probability density function* is defined as

$$p(x) = \frac{1}{x \sigma \sqrt{2\pi}} e^{-\frac{(\log(x) - \mu)^2}{2\sigma^2}}, \quad (6)$$

where μ is the mean and σ is the standard deviation of the logarithm. Based on these variables we are able to represent contact models for situations with varying properties. This enables us to model distinct contact behavior for wiping tasks of arbitrary tool-medium-surface constellations. The log-likelihood based approach allows us to avoid a fixed force threshold by exploiting the variance of the likelihood function. While a fixed force threshold may be sufficient to distinguish contact from no contact, it will often result in false positives in borderline situations. Example probability density function plots are given in Fig. 6. Utilizing a steeply parameterized log-likelihood function (e. g. to simulate a window wiper skimming water from a window, blue in Fig. 6), even path segments with a lower force measurement have a chance to produce the desired outcome. Vice versa, even high contact forces may result in no effect for more flat log-likelihood functions (e. g. to simulate the bristles of a broom collecting fine sand, red in Fig. 6).

The resulting estimation over time for a sponge collecting chippings of glass is illustrated in Fig. 7, where $\mu = 1.0$ and $\sigma = 0.8$. The initial particle distribution is not visually localized. A unified particle distribution is assumed as we aim to solve wiping tasks where no visual feedback is

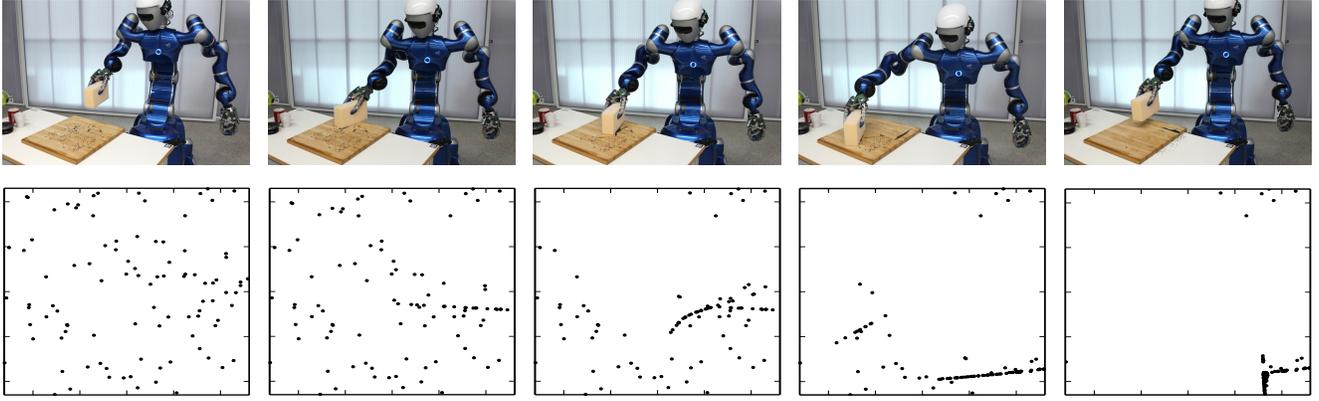


Fig. 7. The real collect action executed by Rollin' Justin (top row) compared to the estimated outcome (bottom row). The chippings are placed to assess the task performance. They are not visually perceived by the robot, instead a uniform initial distribution is assumed.

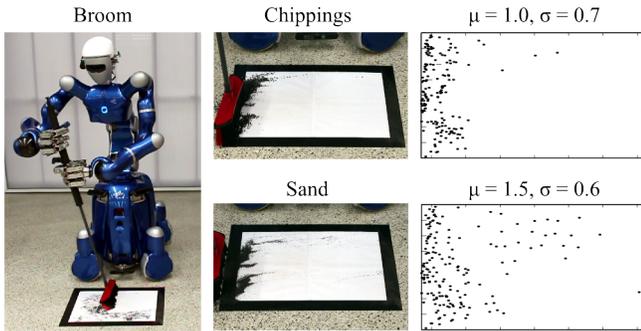


Fig. 8. A broom is used to collect particles (best viewed in a digital copy). The broom is in general less accurate than the sponge utilized in the previous example. Some of the chippings (2 mm - 6 mm) are not effected by the broom (center). The task performance decreases with the size of the particles as it is observed for the fine grained sand (0.1 mm - 2 mm).

available. Since there is only a minor loss of force over a short period of time it is not affecting the estimation too much which reflects the observation of the real execution.

A second scenario showcases the robot using a broom to collect particles on the floor (see Fig. 8). This experiment is executed twice with particles of different size. The robot motion as well as the controller parameterization is identical in both trials. First, the chippings introduced in the previous example are distributed on a sheet of paper on the floor. As the broom swipes over the surface most of the particles are effected. Only a few chippings remain as the bristles of the broom are of irregular nature. Second, small grains of sand are distributed. These particles are too small to be efficiently collected by the broom as the bristles bend. This effect can be modeled with the particle representation by utilizing a flatter log-likelihood function as it is plotted in Fig. 6.

The third scenario investigates the effect of different tools to a certain medium. In particular, we command the robot to remove detergent from a solar panel utilizing the rubber blade of a window wiper versus the bristles of a brush (see Fig. 9). Obviously, the brush is not designed to manipulate liquids. As a result, the window wiper outperforms the brush in this task. To better assess the effect, macro recordings before and after the wiping action have been taken for both

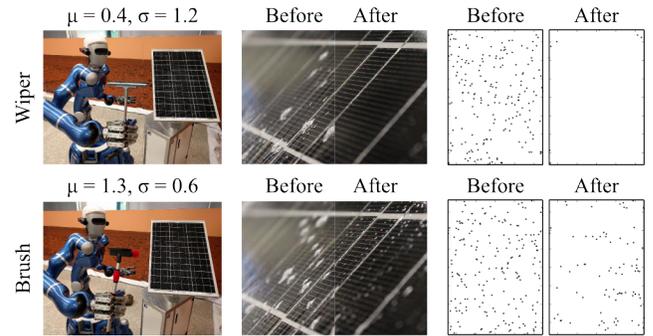


Fig. 9. Detergent is applied to a solar panel. The robot executes a skim action with a window wiper and a brush in order to remove the liquid. The model parameters μ and σ are designed to match the effect of the two tools while the medium stays the same.

trials (center). The window wiper removes most of the liquid from the solar panel surface, whereas a layer of detergent remains for the brush trial. This effect is best viewed in the accompanied video material. Similar to the earlier experiments, this circumstance can be modeled by adapting the parameters of the log-likelihood function. Even though the detergent is hardly visible, the particle model allows the robot to estimate the effect of its actions qualitatively as it is illustrated on the right.

In conclusion, the applied computational model presents a suitable estimation of the real world effect. Even though the contact behavior poses a strong simplification, the resulting patterns in the particle model match the real world observations in general, which allows for a qualitative assessment.

V. FAILURE DETECTION AND RECOVERY

As already emphasized, we utilize an impedance control strategy with a saturated contact force to execute wiping motions with the compliant humanoid robot Rollin' Justin. The utilized control strategy constitutes a safe approach to enable human-robot interaction. However, this makes the approach prone to errors arising from undesired contact situations, i. e. collisions. The effect inference method proposed in this work is capable of detecting these failure states and adapt the effect estimation accordingly.

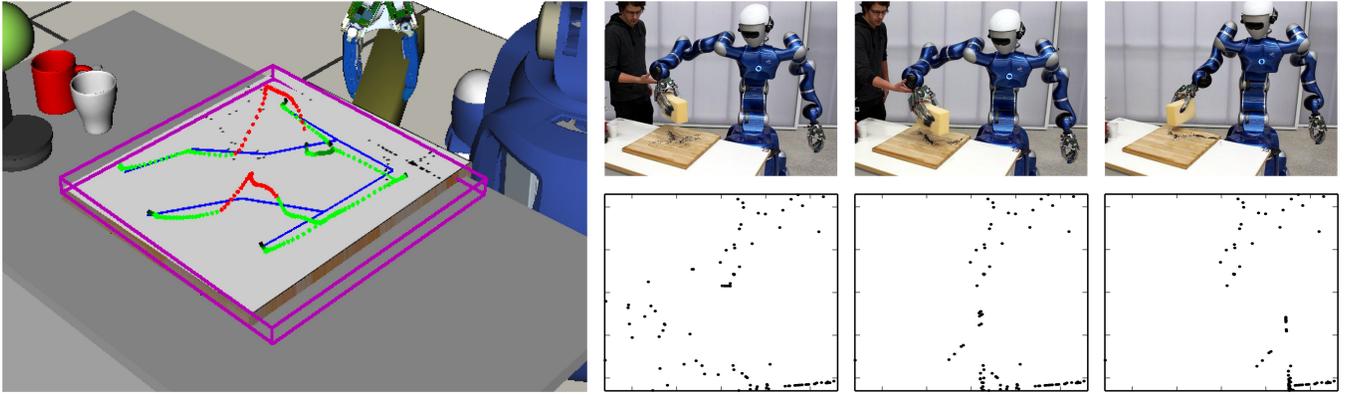


Fig. 10. The Cartesian wiping motion for the failure situation is visualized as red and green dotted path. Red path segments indicate external disturbances which result from the perturbation during task execution as it is shown in the top row. The estimated particle distribution is visualized in the bottom row.

The failure scenario constitutes a deliberate human intervention where the robot is pushed away from the chopping board during the wiping motions. The maximum controller force is saturated to satisfy the desired contact behavior as described in (2). Consequently, the force readings for the human intervention event do not significantly differ from the successful task execution. However, the measured end-effector position deviates from the desired position. The RanSaC algorithm (5) effectively neglects the outliers and returns only the measurements $x_{\text{plane},i}$ that are in contact with the target surface. The outliers are ignored during the effect estimation respectively. The robot is eventually able to correctly infer where the wiping motions have been efficient, and where the desired effect was not carried out. Fig. 10 illustrates the false positive measurements as red dots along the measured Cartesian tool path. The corresponding structured logs for the end-effector motion and the contact

force are provided as overview in Fig. 2. The depicted plot is a two-dimensional visualization of the task execution. The approximated external force $f_{\text{ext},n}$ is plotted in the upper sub-figure and the corresponding TCP height $x_{\text{tcp},z}$ is shown below. The path segments that show a high contact confidence with the chopping board are highlighted in green. The false positive segments that are actually in contact with the human are highlighted in red. Both figures indicate that the robot was pushed twice right before the two main intersection points of the arrow-shaped wiping motion. Three snapshots in the right of Fig. 10 outline the real experiment and the estimated particle distribution respectively. As the robot is disturbed twice at crucial intersection points of the wiping motion, it is not able to accomplish the collect task. A trail of particles is estimated to remain in the central area of the chopping board which is equivalent to the real outcome.

In order to recover from this failure state we apply the reasoning mechanisms described in Sec. III to replan the necessary wiping motions and finish the given cleaning task. The procedure is outlined in Fig. 11. As the robot maintains the described internal effect model there is no need for visual feedback. The robot is able to plan the recovery motion directly in the effect-space of the particle distribution representation projected on the target surface of the chopping board. The KDE is conducted w. r. t. the particle distribution that is estimated after the previously disturbed collect action. Based on the KDE we are able to compute the most significant accumulations of left-over dirt particles (red dots), to be utilized as waypoints for the recovery motion. This recovery motion represents a collect action in itself as it was planned earlier for the initial particle state. Eventually the robot is able to accumulate almost all chippings at the designated goal region.

VI. CONCLUSION

In this work we described a method to infer the effect of wiping motions based on haptic perception rather than visual data. (i) A particle distribution model was utilized to estimate the effect of robotic cleaning actions w. r. t. haptic feedback information, (ii) where a log-likelihood based contact model allows to simulate different tool-medium-surface constella-

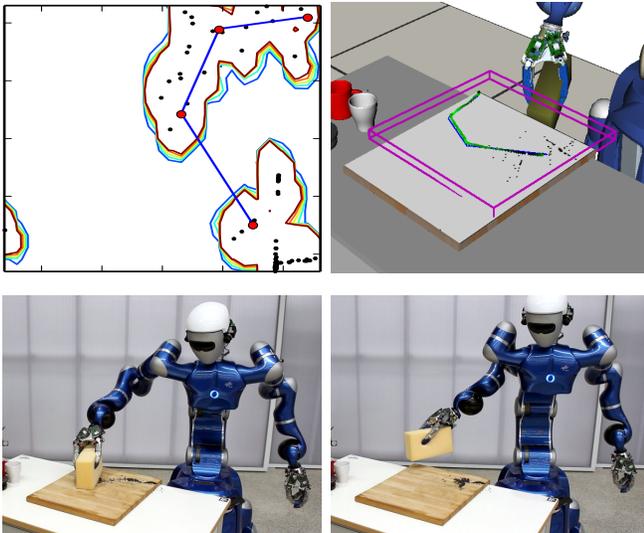


Fig. 11. Top left: The KDE for the remaining particle distribution after the robot was disturbed. The planned tool motion is shown as blue lines. Top right: The recorded recovery motion and the final particle distribution projected on the target surface of the chopping board. Bottom: The real world execution shows the eventually successful task outcome.

tions. (iii) The method is able to infer failure situations due to bad contact arising from human intervention. (iv) We showed that the robot can plan additional wiping motions based on the inferred information in order to successfully accomplish the commanded tasks despite prior failure situations. All experiments are shown in the accompanied video.

Although the results of the effect inference are quite satisfactory, the parameterization of the effect model introduces some limitations. For now, the tool dependent parameters for the log-likelihood function (i. e. μ and σ), as well as the maximum contact force $f_{\text{ext},n,\text{max}}$ were defined empirically. However, Do *et al.* [3] showed that task parameters in the context of wiping motions can be learned over time if visual observation is possible. Given the right parameterization, our method roughly matches the real world outcome for the showcased scenarios. Even though the applied effect model is only a qualitative representation of the actual procedure (i. e. the contact behavior as well as the motion of the particles are strongly simplified), it can be used for a first quality estimation in real-time. A more accurate estimation of the physical behavior could be conducted by a physics simulation such as YadeDEM [22], which is able to simulate the contact behavior between spherical or clustered particles and objects of arbitrary geometric shape in three-dimensional under consideration of physical properties such as friction and internal forces.

An alternative to our approach might be to utilize machine learning algorithms such as random forests or neural networks. Given enough training data, this will enable the robot to identify more sophisticated contact situations as well as more failure cases including abrupt contact loss, permanent human intervention without desired contact at all, and unforeseen impacts as it was shown by Hornung *et al.* [23]. The RanSaC based approach presented in this paper may however provide an initial assumption to apriori label known contact situations (e. g. nominal, tilted, and pushed). As a result the robot might be able to detect bad contact situations already during execution and react online to improve the desired effect of the wiping motion preventive instead of recovering after the motion was performed.

VII. ACKNOWLEDGMENTS

This work was funded by the Helmholtz Association Project HVF-0029 RACE-LAB.

REFERENCES

- [1] J. Hess, J. Sturm, and W. Burgard, "Learning the state transition model to efficiently clean surfaces with mobile manipulation robots," in *Proc. of the Workshop on Manipulation under Uncertainty at the IEEE International Conference on Robotics and Automation (ICRA)*, 2011.
- [2] D. Martínez, G. Alenya, and C. Torras, "Planning robot manipulation to clean planar surfaces," *Engineering Applications of Artificial Intelligence*, vol. 39, pp. 23–32, 2015.
- [3] M. Do, J. Schill, J. Ernesti, and T. Asfour, "Learn to wipe: A case study of structural bootstrapping from sensorimotor experience," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1858–1864.

- [4] M. Kawato, "Internal models for motor control and trajectory planning," *Current opinion in neurobiology*, vol. 9, no. 6, pp. 718–727, 1999.
- [5] J. J. Gibson, "Observations on active touch," *Psychological review*, vol. 69, no. 6, p. 477, 1962.
- [6] J. R. Flanagan, M. C. Bowman, and R. S. Johansson, "Control strategies in object manipulation tasks," *Current opinion in neurobiology*, vol. 16, no. 6, pp. 650–659, 2006.
- [7] G. Hirzinger, N. Sporer, A. Albu-Schäffer, M. Hahnle, R. Krenn, A. Pascucci, and M. Schedl, "Dlr's torque-controlled light weight robot iii-are we reaching the technological limits now?" in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2. IEEE, 2002, pp. 1710–1716.
- [8] D. Leidner, W. Bejjani, A. Albu-Schäffer, and M. Beetz, "Robotic agents representing, reasoning, and executing wiping tasks for daily household chores," in *Proc. of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2016.
- [9] C. Borst, T. Wimböck, F. Schmidt, M. Fuchs, B. Brunner, F. Zacharias, P. R. Giordano, R. Konietschke, W. Sepp, S. Fuchs *et al.*, "Rollin'justin-mobile platform with variable base," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009, pp. 1597–1598.
- [10] A. Gams, T. Petric, B. Nemeč, and A. Ude, "Learning and adaptation of periodic motion primitives based on force feedback and human coaching interaction," in *IEEE-RAS International Conference on Humanoid Robots (ICHR)*, 2014.
- [11] J. Hess, M. Beinhofer, and W. Burgard, "A probabilistic approach to high-confidence cleaning guarantees for low-cost cleaning robots," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 5600–5605.
- [12] M. Ghallab, A. Howe, D. Christianson, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins, "Pddl—the planning domain definition language," *AIPS98 planning committee*, vol. 78, no. 4, pp. 1–27, 1998.
- [13] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning: theory and practice*. Morgan Kaufmann, 2004.
- [14] L. Kunze, M. E. Dolha, E. Guzman, and M. Beetz, "Simulation-based temporal projection of everyday robot object manipulation," in *Proc. of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2011, pp. 107–114.
- [15] J. Winkler and M. Beetz, "Robot action plans that form and maintain expectations," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 5174–5180.
- [16] P. Pastor, M. Kalakrishnan, S. Chitta, E. Theodorou, and S. Schaal, "Skill learning and task outcome prediction for manipulation," in *Robotics and Automation (ICRA)*, 2011 *IEEE International Conference on*. IEEE, 2011, pp. 3828–3834.
- [17] D. Leidner, C. Borst, A. Dietrich, and A. Albu-Schäffer, "Classifying compliant manipulation tasks for automated planning in robotics," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 1769–1776.
- [18] A. Dietrich, T. Wimböck, A. Albu-Schäffer, and G. Hirzinger, "Reactive whole-body control: Dynamic mobile manipulation using a large number of actuated degrees of freedom," *IEEE Robotics & Automation Magazine*, vol. 19, no. 2, pp. 20–33, 2012.
- [19] D. Leidner, A. Dietrich, M. Beetz, and A. Albu-Schäffer, "Knowledge-enabled parameterization of whole-body control strategies for force-sensitive tasks," *cond. accepted for publication in Autonomous Robots (AURO): Special Issue on Whole-Body Control of Contacts and Dynamics for Humanoid Robots*, 2015.
- [20] S. Denei, F. Mastrogiovanni, and G. Cannata, "Towards the creation of tactile maps for robots and their use in robot contact motion control," *Robotics and Autonomous Systems*, vol. 63, pp. 293–308, 2015.
- [21] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [22] V. S. *et al.*, "Yade documentation," *The Yade Project*, vol. 2nd ed, 2015.
- [23] R. Hornung, H. Urbanek, J. Klodmann, C. Osendorfer, and P. Van Der Smagt, "Model-free robot anomaly detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2014, pp. 3676–3683.