

DOF-Decoupled Active Force Sensing (D-DAFS): A Human-Inspired Approach to Touch-based Localisation Tasks

Niccoló Tosi¹, Olivier David¹, Herman Bruyninckx²

Abstract—In the context of touch-based object localisation in structured environments, solving the full 6D problem is computationally expensive as complexity scales exponentially with the number of degrees of freedom. Previous related works provided solutions to problems whose initial uncertainty was bounded by the computational feasibility for on-line applications. Observing the results of an experiment in which human beings face a blind localisation task, a new planning approach, named DOF-Decoupled Active Force Sensing (D-DAFS), is presented for such robotics applications. Specifically, the whole task is divided into a sequence of subtasks, each of them focused on reducing a piece of uncertainty, so that the resolution of the inference model can be increased accordingly. This allows a better allocation of resources, focusing on adopting the most accurate, thus expensive, scheme only when the uncertainty is sufficiently low for on-line calculations. The proposed methodology has been applied to perform a 3-DOF localisation with a Staubli RX90 robot, and a series of simulations have been run to prove its effectiveness in coping with high initial uncertainty avoiding computational overheads.

I. INTRODUCTION

In the context of scene-model calibration, force sensing widens the set of scenarios in which object localisation is made possible, both as a stand-alone tool and as part of a multi-sensor system. Although a good number of works have studied the problem of processing force or tactile information to infer object pose, current touch-based localisation techniques are limited to tasks in which the initial uncertainty is in the order of 0.5m. More specifically, as the solution space is increased, the problem becomes less and less computationally tractable using non-linear Bayesian inference schemes, e.g. grid-based or particle-filter methods, since complexity scales exponentially with the number of DOFs considered.

In order to cope with higher initial uncertainty, this paper introduces a new, computationally-cheap approach, focused on reducing the complexity of the task and on adapting the level of resolution of the estimator as the localisation takes place. Inspired by the observation of a blind localisation experiment



Fig. 1. A subject during the localisation test.

carried out with 30 subjects (Figure 1), the proposed approach is named DOF-Decoupled Active Force Sensing (D-DAFS).

Performing scene-model calibration with efficiency is a major concern for tele-operation applications such as those depicted in [1]. In particular, in case of harsh environments it is necessary to cope with highly-soiled scenes, e.g. by dust, mud and dirty water. Figure 2 shows an example of real-life scenario in which the human is responsible for the gross motion strategy, while the robot fills in the fine motion and does the calibration. For such applications, force sensing is preferred over tactile sensing as its performance is less dependent on the finesse of the touched surface and the scale of the interaction forces are in the range of 1000N.




Fig. 2. Example of harsh environment in which calibration is necessary

This paper is organised as follows: Section II summarises the state of the art on touch-based localisation, Section III describes the experiment carried out with human subjects, Section IV illustrates the proposed methodology and Section V provides the results obtained applying D-DAFS to a box-localisation task.

*This work was supported by CEA, LIST, Interactive Robotics Laboratory.

¹N. Tosi and O. David are with CEA, LIST, Interactive Robotics Laboratory, Fontenay aux Roses - F92265, France {niccolo.tosi}, {odavid}@cea.fr

²H. Bruyninckx is professor at the Department of Mechanical Engineering, University of Leuven, Belgium herman.bruyninckx@mech.kuleuven.be

II. STATE OF THE ART

A. Touch-Based Localisation

In the early 80s, 3D and 6D polyhedral-object localisation was performed using interpretation trees [2], [3]. Then, probabilistic schemes were developed to accomplish the localisation accounting for geometric and measurement uncertainties. Though powerful, probabilistic approaches tend to become computationally expensive as the problem complexity increases with the initial uncertainty and the number of DOFs. In order to limit the size of the problem, previous works analysed touch-based localisation reducing either or both the number of DOFs and the search space dimensions.

In 2001, Gadeyne et al. [4] performed cube localisation in 3DOFs by sampling the solution space and processing force-sensor information. In 2005, Gadeyne et al. [5] adopted sequential Monte Carlo sampling to estimate contact-state formations during cube-in-corner operations using force sensing. In 2006, Petrovskaya et al. [6] performed 6DOF force-based localisation with uncertainty of 400mm in position and 360 deg in orientation, proposing an enhanced Particle Filter scheme named Scaling Series method in a later paper [7].

B. Active Sensing

In the framework of robotic localisation tasks, the definition of an optimised sequence of actions aimed to reduce the uncertainty over the pose of the environment elements goes under the name of active sensing.

Previous related works analysed the problem of actively sensing the environment for 3D self-localisation tasks. In [8], Markov localisation was applied to infer robot position inside a structured office environment using laser and ultra-sound sensors. A cost-vs-utility function was introduced to select the best movements and sensing actions to perform. Costs were represented by time and energy spent, and utility was represented by the expected decrease of uncertainty evaluated through the Kullback Leibler divergence (see [9]), i.e. the non-symmetric measure of the difference between the current posterior probability distribution and the envisioned Probability Density Function (PDF) after the sensing action. Further to that work and focusing on reducing the sensor resources required, the same task was performed using angular and linear odometers, a compass and a contact sensor, an angular odometer and a contact sensor [10]. More recently, the same operation was accomplished only using a clock and a contact sensor [11].

Active force sensing to perform cube-in-corner tasks with autonomous motions was studied in [12]. The full operation was decoupled into a compliant sequence of contact formations (CFs), stored in the so-called CF Graph. Then, the optimisation process was decoupled into a requirement for the CF sequence and a requirement for the active sensing motions in each CF. Results from this work proved that such an exhaustive approach is computationally expensive and suggested further research to adjust the level of resolution of the estimator and controller to the accuracy required by the single action.

In 2009, Hsiao [13] performed object localisation and grasping in 3DOFs using tactile sensors installed on a three-finger hand. The authors focused on the decision-making side of the problem by implementing an approximated Partially Observable Markov Decision Process (POMDP) [14] scheme to maximise the information gain over a finite-horizon series of motions. The best action is chosen by using the Kullback Leibler divergence as a metric to decide among a pre-determined set of trajectories, defined relatively to the best pose estimate of the object. The measurement model accounts for both contact and no-contact information.

In 2012, further to Hsiao's work, Hebert et al. [15] performed 3DOF localisation using a hand with force sensors. With respect to Hsiao, motion primitives are generated and evaluated as the localisation takes place. Furthermore, Javdani et al. [16] compared the information-based approach followed by Hebert and Burdick with a hypothesis-pruning technique through a series of simulations.

C. Discussion

State-of-the-art techniques of touch-based localisation are facing two major challenges: performing tasks with higher initial uncertainty, and finding a solution for the problem of where-to-sense-next. In this paper, we focus on the former. To our best knowledge, one of the major limits of nowadays active force sensing lies in the use of maximum resolution at both estimator and control level throughout the whole localisation process, independently from the current uncertainty and task objectives.

III. HUMAN APPROACH TO BLIND LOCALISATION

Driven by the need to analyse the decision-making process occurring during touch-based localisation operations, a test was carried out to observe human approach to blind localisation tasks. The aim of the experiment was to observe human beings performing a typical robot localisation task using a force sensor. The test was designed to reproduce the uncertainty and measurement system occurring on an anthropomorphic robot arm featuring a spherical end-effector with a wrist force sensor. To simulate such conditions, subjects were asked to use a 20-cm stick to perform the localisation. Other forms of scene exploration were forbidden. To simulate robot measuring system conditions, it was necessary to shut down all the senses apart from wrist force sensing. Therefore, all the subjects were asked to (see also Figure 1):

- wear thick gloves to prevent finger tactile sensing
- wear an eye-cover to prevent vision
- wear a headset to prevent hearing.

The test focused on understanding whether:

- 1) the subjects performed the localisation task by reducing the uncertainty DOF-by-DOF
- 2) common patterns could be identified among subjects' strategy choice.

A. Experiment Set Up

The test was carried out at CEA, LIST, and was taken by 30 adult subjects. A schematic representation of the experiment scene is illustrated in Figure 3. The elements present in the scene were:

- two protections: $p1$ and $p2$
- a square table $T_{\{ABCD\}}$ with side of 1m
- the V-block V shown in Figure 3

Our analysis assumes that the localisation task was performed starting with uncertainty on the DOFs presented in Table I.

DOF	Description
γ_P	tester orientation around the z axis
x_T	table x coordinate
y_T	table y coordinate
γ_T	table orientation around the z axis
x_{VB}	V-block x coordinates
y_{VB}	V-block y coordinates
γ_{VB}	V-block orientation around the z axis

TABLE I
SYSTEM DOFS




Fig. 3. Experiment set up and V-block

Particular attention was paid to prevent each subject from knowing the scene before her/his trial started. It was also forbidden to attend other people's trials. As the test started, the subjects had 30s to accurately observe all the elements composing the scene before wearing the glove, the eye cover and the headset. This way, they could hold a 3D representation of the environment while performing the blind localisation. The experiment was recorded using a camera positioned beyond the table with respect to the tester. A few samples of the recorded videos are available online [17].

Two trials were carried out to let the subjects face two different uncertainty situations, referred to as Test 1 and Test 2.

- **Test 1 - localisation of the V-block on the table:** Subjects were given the information that the object was located on the table. Position and orientation of the object were unknown.
- **Test 2 - localisation of the V-block on the table edge:** Subjects were given the information that the object was

located on the AB edge and oriented with the V-valley parallel to AB edge.

In both tests, subjects were asked to position themselves on the circle S , spin once, and then localise the V-block. The task was considered accomplished when the stick was steadily released in a horizontal position inside the V-block vertical valley in less than 120s.

B. Experiment Results

From the experiment observation, the level of initial uncertainty experienced by the subjects changed of intensity from person to person, due to the sense cutting-off and the initial spinning. We consider this mainly to depend on: self-orientation skills, fear of darkness and body equilibrium.

While performing the task, the subjects followed common behavior patterns which can be divided into a series of actions directly focused on reducing the uncertainty over a set of DOFs, as shown in Table II. For each step, Table II presents the identified action and its correlation with the system uncertainties expressed as DOFs. For both tests, the number of subjects who decided to take the action is reported.

Step	Chosen Action	DOF(s)	Test 1	Test 2
1	touch $p1$ or $p2$	γ_P	6/30	5/30
2	touch AC	γ_T x_T	26/30	25/30
3	touch AB or CD	y_T	16/30	21/30
4	table-compliant	z_{VB}	21/30	22/30
5	T exploration	x_{VB} y_{VB}	30/30	30/30
6	V exploration	γ_{VB}	30/30	30/30

TABLE II
CORRELATION OF DOFS AND ACTIONS

Although not all of the testers followed every step of the series, the sequence never changed. In other words, a few experimenters skipped some steps, depending on their prior knowledge and confidence about the specific DOF, but always followed the same action order. Since all the subjects were given the same information and freely chose which strategy to follow, the link between the uncertainty and the sequence of actions is to be found in the topological and dynamic relationships between scene objects. More specifically, the knowledge of the V-block being on the table was interpreted by the subjects who decided to approach the problem by

- 1) localising the protections
- 2) localising the table
- 3) localising the V-block on the table.

After step 4 was performed and a sufficient knowledge of the table pose was achieved, the subjects faced the problem of exploring the table surface to locate the V-block. Unlike previous steps which were skipped by some subjects, nobody missed steps 5 and 6 during both Test 1 and 2.

Having examined the results of the experiment, the following conclusions can be made:

- Human beings approached the proposed localisation problem by dividing it into a sequence of steps, so a feasible task breakdown could be identified.

- A direct correlation can be established between the actions performed by the subjects in each step and the tackled DOFs.
- All the subjects processed the same available information and freely chose to follow the same order of actions.

IV. DOF DECOUPLING

A. Overview

Our research is driven by the need to develop robotic localisation strategies able to adapt the level of resolution on the estimator, controller and sensors with respect to each operation involved in the task. In this regard, our aim is to merge the minimalistic-sensing strategy adopted in [8] [10] [11] with the step-by-step approach followed in [12]. We provide a new scheme for n-dimensional localisation tasks, applicable to multi-sensor systems and able to reduce the resources required in the operation. Further to the results of the experiment presented in Section III, we focus on decoupling the problem of active sensing into a sequence of exploring actions. Each of them is aimed to reduce a portion of the global uncertainty. For each action, the parameters of the estimator and controller can be set so that they are consistent with the required level of resolution.

In particular, we aim to:

- adapt the level of resolution to be consistent with the current uncertainty
- adopt different model representations of the object as the localisation advances
- formalise the decision-making process as a trade-off accounting for information-gain, motion and calculation

Without explicit formalisation, a similar approach was introduced in an earlier work [18], tackling the problem of localising a hook in 3D for submarine applications. Previous work introduced action-decomposition for POMDPs [19]. With respect to that approach, in our formulation the problem is decomposed through the configuration DOFs, and the level of resolution is adapted throughout the localisation process.

B. DOF-by-DOF Problem Decoupling

To specify the localisation *task* with DOF Decoupling, the following must be defined:

- uncertainty parameters $\mathbf{x} = \{x_i\}$ defining the pose of the object with respect to a reference frame, representing the DOFs of the localisation task
- initial uncertainty values for \mathbf{x}
- available robot actions and control modes, e.g. guarded point-to-point motion or compliant motion
- available sensors, e.g. contact detector, force-torque sensor, laser scanner or camera.

Previous related work proved the full 6D action planning optimisation not feasible for online applications [12]. Here, the initial set of uncertainty DOFs is to be divided into a sequence of subsets, so that the localisation procedure can be decomposed into *subtasks* that are easier to solve (*divide et impera* approach). In case of large initial uncertainty, this

allows to start with a coarse geometric model and low-weight estimator when the information is poor, and to progressively refine them as more measurements are available to infer the object pose with more accuracy.

Each *subtask* is defined by:

- the target DOF subset - the set of degrees of freedom whose uncertainty is to be reduced in the subtask
- the required maximum level of uncertainty on the DOF subset
- the adopted robot motion mode
- the geometric model used to represent the objects in the scene

Figure 4 graphically represents the high-level subtask decoupling scheme described above.




Fig. 4. Subtasks decoupling scheme

Within each *subtask*, robot actions are chosen to maximise the collected information to reduce the uncertainty on the current DOF subset, also accounting for the cost of motion and computation. This problem is formalised as a greedy POMDP, in order to keep it tractable for on-line applications. The reward function r accounts for the expected information gain I , the motion cost C_m and the computational cost C_c of a given action a executed under belief state b . The weight factors $\{w_I, w_m, w_c\}$ give different impact to the aforementioned metrics.

$$r(b, a) = w_I I(b, a) - w_m C_m(b, a) - w_c C_c(b, a) \quad (1)$$

Although this paper focuses on touch-based localisation, Eq. 1 can be also applied to make decisions about different forms of sensing, evaluating their cost and benefit with respect to the current uncertainty. In some applications, a single type of sensor is efficient only in a part of the global localisation process. As for the calculation of the expected information I , in [20], methods based on covariance matrix for unimodal distributions are presented, together with methods for multi-modal distributions such as entropy, reduction of entropy and Kullback-Leibler Divergence (KLD). Here, I is defined as the KLD between the current posterior P_t and the posterior probability P_{t+1} after executing action a , under state belief b .

$$D_{KL}(P_t || P_{t+1}) = \int P_{t+1} \log \frac{P_{t+1}}{P_t} dx \quad (2)$$

With a time horizon $T = 1$, the chosen action a^* is the one that maximises the immediate reward.

$$a^* = \arg \max_a (r(b, a)) \quad (3)$$

As for the motion cost metric C_m , suitable options may be represented by the energy spent or the task-execution time, normalised with respect to the energy and time of some reference action. C_c can be calculated by carrying out computation trials to measure the cost of processing the information as number of elementary operations or execution time.

V. RECTANGLE LOCALISATION

Here, we present an application of DOF-Decoupled Active Sensing to a 3D localisation task in which the object is fixed. With respect to the human experiment, we focus on steps 5 and 6 of the sequence, localising a solid rectangle positioned on a table with high initial uncertainty (up to 2m in translation). Such task would become not feasible online if treated with tradition methods, i.e. with full resolution from the beginning. With respect to the v-block used in the experiment, the simpler rectangular shape allows to calculate the likelihood function analytically. Nevertheless, the decoupling scheme is rather generic, and can be applied to objects of higher complexity.

A. Task Description

The task is to localise a rectangle Q which can assume any $\{x_Q, y_Q\}$ position on top of a table, and can have any orientation γ_Q around the z axis. We suppose the robot to be equipped with a spherical end effector coupled with a force-torque sensor. The robot has to localise Q exploring the environment by sweeping with the end effector.

B. Task Decoupling

Further to the methodology introduced in Section IV, the task is decoupled into two subtasks: $S1$ and $S2$. The applied scheme, i.e. which DOFs to tackle in which subtask, is taken as-is from the one empirically observed during the experiment described in Section III. Specifically, $S1$ is aimed to reduce the uncertainty on x_Q and y_Q , with $S2$ focusing on refining the estimation on x_Q , y_Q and γ_Q . To simplify the estimation, Q is represented as its bounding circle during $S1$, whereas it is modelled as a rectangle in $S2$. The first contact with the object is used as triggering event to switch from $S1$ to $S2$.

C. Subtask $S1$

- Target DOFs: $\{x_Q, y_Q\}$
- Geometric model: bounding circle of radius ρ
- Maximum level of uncertainty: linear uncertainty bounded by ρ
- Motion mode: point-to-point guarded motion

1) *Inference scheme*: During $S1$, the solution space is represented by the whole table. At this level of resolution, the object is represented as a circle, and the posterior probability $P(x_Q, y_Q | \mathbf{m})$ is built over a uniform grid covering the solution space. To guarantee a sufficient level of resolution, the maximum distance between grid nodes is set smaller than the short side of Q . Each node corresponds to the

centre of the bounding circle that approximates one of the possible configurations of the object. Since we assume no prior knowledge about x_Q and y_Q , $P(x_Q, y_Q)$ is initialised as a uniform distribution.

2) *Measurement model*: The adopted measurement model takes into account both negative and positive information. This means that both contact and no-contact information are used to enhance the localisation, and the sweeping performed before the actual contact is part of the localisation process. Practically, the swept area is defined by the path τ followed by the end effector's reference frame and the thickness ξ of the sweep. S_τ and F_τ represent the initial and final point of the sweep path τ , as depicted in Figure 5.




Fig. 5. Sweep measurement models for subtasks $S1$ and $S2$

There are two possible readings from our contact-detection sensor, namely "no contact" and "contact".

$$\mathbf{m} = \{\text{NC}, \text{C}\} \quad (4)$$

Intuitively, when a sweep intersects the circle surface corresponding to one of the grid nodes, the likelihood of measuring NC or C is a function of the intersected area. Similarly to the measurement model used in [13], in this first, low-resolution subtask, we simplify the calculation by using the deepest collision distance d as metric for such intersection.

$$P(\mathbf{m} | x, y) = P(\mathbf{m} | d) \quad (5)$$

Although some information is lost, this method is computationally cheap and allows to quickly discard low-probability configurations which are swept through. Specifically:

$$P(\mathbf{m} = \text{NC} | d) = 1 - \alpha(d) \quad (6)$$

$$P(\mathbf{m} = \text{C} | d) = \alpha(d) \quad (7)$$

The definition of $\alpha(d)$ depends on the geometric dimensions of the bounded object. In our case, this function is calculated as the likelihood of measuring contact with the rectangular object bounded by the circle, after marginalisation over the rotation γ around the z axis.

$$P(\mathbf{m} = \text{C} | d) = \int_0^{2\pi} P(\mathbf{m} = \text{C} | d, \gamma) d\gamma \quad (8)$$

Eq. 8 can be solved off-line to build a look-up table for different values of d and speed up the run-time execution. In our implementation, the calculation is approximated by a

simple model where the parameter α is a linear function in d which is bounded: so $d \in [0, \rho]$, thus $\alpha \in [0, 1]$.

$$\alpha(d) = \frac{d}{\rho} \quad (9)$$

In practice, in case of a sweep with no contact, Eq. 6 is used to calculate the likelihood of each grid node. In case of a sweep ended with a contact, this is treated as a no-contact sweep, applying Eq. 6, plus an infinitesimal contact sweep, applying Eq. 7. A similar approach was introduced in [13].

3) *Decision Making*: During $S1$, the decision-making process about where-to-sense next is simplified by executing a pre-programmed spanning movement over the table, parametrised with respect to the shortest side of the rectangle, with increments on the x axis equal to the node distance. This guarantees an exhaustive exploration of the table similar to the one observed during the experiment described in Section III. Figure 6 shows a screenshot while subtask $S1$ is running.




Fig. 6. Parametric exploration strategy. The swept path is coloured in green, the pose of the rectangle in red, and the circumferences corresponding to the considered $\{x_Q, y_Q\}$ configurations of the object are depicted in blue.

D. Subtask $S2$

In the framework described in Section IV, subtask $S2$ is defined by:

- Target DOFs: $\{x_Q, y_Q, \gamma_Q\}$
- Geometric model: rectangle
- Maximum level of uncertainty: linear error smaller than 5mm, orientation error smaller than 5 deg
- Motion mode: point-to-point guarded motion

After the first contact, the transition to Subtask $S2$ is triggered. The geometric model is refined substituting the bounding circle with a rectangle.

1) *Inference scheme*: In $S2$, we adopt a Particle Filter (PF) inference model to account for the multi-modal posterior distribution over the pose of the object. The target DOF subset includes $\{x_Q, y_Q, \gamma_Q\}$. Three-dimensional samples are drawn from the posterior distribution over $\{x_Q, y_Q\}$ inherited from $S1$, and a uniform distribution over γ_Q . After each sensing action, the posterior is updated and zero-probability particles are pruned.

2) *Measurement model*: Similarly to $S1$, both positive and negative information are taken into account, and the sweep τ is parametrised as depicted in Figure 5. Though, in this case, both contact and no-contact likelihoods are defined with respect to

the intersected area A_{int} , instead of the contact depth d (see Figure 5). Formally:

$$P(\mathbf{m}|x, y, \gamma) = P(\mathbf{m}|A_{int}) \quad (10)$$

$$P(\mathbf{m} = \text{NC}|A_{int}) = 1 - \alpha \quad (11)$$

$$P(\mathbf{m} = \text{C}|A_{int}) = \alpha \quad (12)$$

The α function is defined as a linear function of A_{int} .

$$\alpha = \frac{A_{int}}{A_{max}} \quad (13)$$

A_{max} is the maximum sweepable area of the object, which may or may not correspond to the total surface area of the rectangle, depending on the dimension of the sweep. Using A_{int} to model the contact detection sensor allows to better evaluate to what extent the sweep covers the surface of a single particle, and represents an improvement of the method proposed in [13]. This approach can be extended to model 3D sweeps, introducing the concept of intersected volume.

From a practical point of view, the sweep with no contact is treated as in $S1$, but using the likelihood functions in Eq. 11 and Eq. 12. As the contact is established, with a force-torque sensor able to measure a contact point p_C and a normal vector n_C , previous works in literature [6] [21] calculated the contact likelihood assuming the two measurements to be independent.

$$P(\mathbf{m}|x, y, \gamma) = P(p_C|x, y, \gamma)P(n_C|x, y, \gamma) \quad (14)$$

If only p_C is available, $P(n_C|x, y, \gamma)$ is discarded from Eq. 14. For a single face of each particle, the contact likelihood is calculated as in Eq. 15, with Δ_{fc} being the measured distance between face f and p_C , and σ_P the standard deviation error on the contact point measurement. Similarly for the normal vector, with Δ_{fn} and σ_n as angular distance and standard deviation error. To solve the correspondence problem, i.e. finding which face to consider for each particle, a maximum-likelihood approach is followed here. Specifically, for each particle, the face that is closest to the measured point is the one that "represents" it (Eq. 15):

$$P(\mathbf{m} = \text{C}|x, y, \gamma) = \max_f \{ \mathcal{N}(\Delta_{fc}, 0, \sigma_P^2) \mathcal{N}(\Delta_{fn}, 0, \sigma_n^2) \}. \quad (15)$$

Such approach is sometimes referred to as hard assign.

3) *Decision Making*: In this subtask, a pre-programmed open-loop strategy defined relatively to the updated best estimate is compared to an info-gathering one, as explained below. The former can only be generated for simple geometries, whereas the latter can be extended to arbitrary complex objects.

- **Open-loop strategy**: sensing actions are defined with respect to the current best estimate which is updated after each touch, aiming at the four faces of the object in a pre-defined sequence. Such strategy does not guarantee to fully constrain the localisation problem as, in theory,

the robot could end up touching the same face multiple times.

- **Info-driven strategy:** After each contact, sensing actions are chosen to maximise a simplified reward function which only accounts for the information gain:

$$r(b, a) = I(b, a). \quad (16)$$

The set of actions to choose from is the same as in the previous case (Figure 7), but the sequence is not pre-determined. After three contacts, the sensing is considered finished.




Fig. 7. Open-loop and info-driven decision-making relatively to the updated best estimate (coloured in blue)

E. Simulation Results

A series of simulations were carried out in Matlab to validate the proposed methodology facing a problem with high initial uncertainty. Initially, a 3DOF touch-based localisation task of a 0.25m x 0.15m rectangle was performed with 1m uncertainty in translation and 180 deg in orientation with a contact estimator with $\sigma_P = 0.005\text{m}$, following the open-loop decision making approach in S2. Trials with 2000, 4000, 8000 and 16000 particles for Subtask S2 were performed in order to measure the effect of this parameter on the estimation. Figure 8 illustrates the recorded error on the final estimation after four sensing actions in S2. For this specific case, data show that 16000 particles allow to reduce the error due to poor sampling within the contact sensor standard deviation. Few failures were still recorded due to the fact that the chosen motion strategy could not guarantee to establish contact on each of the four sweeps, as observed in the robot trials. To better explain the benefits of D-DAFS, Figure 9 presents the results obtained simulating the same type of localisation with initial uncertainty of 2m in translation and 180 deg in orientation. The computational time difference between the two trials was negligible, even if the initial search space was increased by a factor of 4. Specifically, the overall computational time to process the information was the same as the one required by the 1m x 1m application (processing each sweep took approximately 0.2s in both trials), since the increased initial uncertainty was entirely tackled in the low-weight Subtask S1.

Considering the higher initial uncertainties with respect to literature examples [7] [13], the achieved result represents a step forward towards on-line applications.




Fig. 8. Final estimation error - linear uncertainty: 1m




Fig. 9. Final estimation error - linear uncertainty: 2m

F. Robot Implementation

The exploration scheme introduced in the previous sections was applied to localise a solid-rectangle using a Staubli RX90 robot equipped with a force-torque sensor coupled with a spherical end effector, as shown in Figure 10. Controlled with a TAO2000 V2 platform [22] running on a VxWorks machine, the robot performed guarded motions with force-based contact detection. From previous tests, we set the contact threshold to 50N, $\sigma_P = 0.005\text{m}$ and $\sigma_n = 6\text{ deg}$. The task was to identify the pose of the object (dimensions: 0.18m x 0.1m x 0.1m) which was fixed on top of a table. The initial uncertainties on x_Q, y_Q, γ_Q were 0.5m, 0.8m and 180 deg respectively. The global task was decoupled into S1 and S2, as described in Section V. In S2, the number of samples used in the particle filter was set to 16000, further to the simulation results. The task was considered accomplished if the end effector could be placed into one of the two up-facing bottle tops with diameter of 30mm, fixed at symmetric known locations, as depicted in picture no. 4 in Figure 10. With the diameter of the end effector equal to 25mm, the allowed maximum linear error was 5mm, which is comparable to state-of-the-art examples [7]. Due to the symmetry of the object, two shots were allowed to execute the peg-in-hole operation.

Tests were carried out with the object assuming 9 different poses defined by 3 random $\{x_Q, y_Q\}$ positions and 3 γ_Q orientations. A video demo is available online [23]. Three runs were performed for each of the 9 poses, following both the open-loop and the info-driven strategy, respectively obtaining a success rate of 23/27 and 25/27. In the former case, failures occurred because the strategy caused the robot to




Fig. 10. 1: table exploration ($S1$). 2: first contact. 3: one of the info-gathering actions performed during object exploration ($S2$). 4: task accomplished.

touch the same face several times, thus not providing sufficient information. Instead, the two failures observed in the latter case were due to a bad FT sensor reading caused by an unmodelled contact formation on a corner.

VI. CONCLUSIONS

This work has introduced D-DAFS as a new approach to touch-based localisation tasks, improving the computational tractability with respect to traditional methods: the full task is decoupled into a series of subtasks, each reducing a portion of the global uncertainty, stepping from low-resolution configuration space to eventually full resolution configuration space (typically 6DOFs).

This *divide et impera* approach is inspired by an experiment carried with thirty human subjects facing a blind localisation task, in which the challenge was to find the table (2DOFs configuration space), the object on the table (3DOFs) and finally the geometric features of the object (6DOFs) [17].

The middle part of this strategy was transferred to a 6-axis manipulator to perform a 3DOF touch localisation task, implemented with a force-sensor-based guarded motion. The achieved localisation started with an initial uncertainty of the table size (0.5m x 0.8m) and allowed the robot to position itself within 5mm from a designated feature on the object with a success rate of 25 out of 27 trials following an info-driven strategy to decide where to sense next.

In addition, the simulation has proven the computational effectiveness of the proposed methodology in coping with problems with initial uncertainty about a magnitude order higher than state-of-the-art examples, simulating the same task with table sizes of 1m x 1m and 2m x 2m. The computational cost scaled better than linearly with table size, significantly improving its on-line feasibility.

VII. ACKNOWLEDGMENTS

The authors wish to acknowledge the support by KU Leuven Geconcerteerde Onderzoeks-Actie *Global real-time optimal control of autonomous robots and mechatronic systems*, and the European Union's FP7 projects *BRICS* (231940), *ROSETTA* (230902), and *RoboHow.Cog* (288533).

REFERENCES

- [1] P. Garrec, F. Geffard, O. David, F. Russotto, Y. Measson, and Y. Perrot, "Telerobotics research and development at cea, list," in *ANS EPRRSD - 13th Robotics & remote Systems for Hazardous Environments . 11th Emergency Preparedness & Response Knoxville, TN*, 2011.
- [2] P. C. Gaston and T. Lozano-Perez, "Tactile recognition and localization using object models: The case of polyhedra on a plane," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 6, pp. 257–265, 1983.
- [3] T. Lozano-Perez and W. E. L. Grimson, "Model-based recognition and localization from sparse range or tactile data," 1983.
- [4] K. Gadeyne and H. Bruyninckx, "Markov techniques for object localization with force-controlled robots," 2001.
- [5] K. Gadeyne, T. Lefebvre, K. Gadeyne, T. Lefebvre, and H. Bruyninckx, "Bayesian hybrid model-state estimation applied to simultaneous contact formation detection and geometrical parameter estimation," *Int. J. Robotics Research*, vol. 24, pp. 615–630, 2005.
- [6] A. Petrovskaya, O. Khatib, S. Thrun, and A. Y. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *In Proc. of ICRA*, 2006.
- [7] A. Petrovskaya, "Towards dependable perception: Guaranteed inference for global localization," 2010.
- [8] D. Fox, W. Burgard, and S. Thrun, "Active markov localization for mobile robots," *Robotics and Autonomous Systems*, vol. 25, pp. 195–207, 1998.
- [9] S. Kullback and R. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.
- [10] J. M. O'Kane and S. M. LaValle, "Localization with limited sensing," *IEEE Transactions on Robotics*, vol. 23, pp. 704–716, August 2007.
- [11] L. Erickson, J. Knuth, J. M. O'Kane, and S. M. LaValle, "Probabilistic localization with a blind robot," in *Proc. IEEE International Conference on Robotics and Automation*, 2008.
- [12] T. Lefebvre, H. Bruyninckx, and J. De Schutter, "Task planning with active sensing for autonomous compliant motion," *The International Journal of Robotics Research*, vol. 24, no. 1, p. 61, 2005.
- [13] K. Hsiao, M. I. of Technology, Dept. of Electrical Engineering, and C. Science, *Relatively Robust Grasping*. Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2009. [Online]. Available: <http://books.google.fr/books?id=tT0SAAACAAJ>
- [14] A. R. Cassandra, L. P. Kaelbling, and M. L. Littman, "Acting optimally in partially observable stochastic domains," pp. 1023–1028, 1994.
- [15] P. Hebert, J. Burdick, T. Howard, N. Hudson, and J. Ma, "Action inference: The next best touch," *Robotics: Science and Systems Workshop on Mobile Manipulation*, 2012.
- [16] S. Javdani, M. Klingensmith, D. Bagnell, N. Pollard, and S. Srinivasa, "Efficient Touch Based Localization through Submodularity," *ArXiv e-prints*, Aug. 2012.
- [17] Video, "Human subjects performing a blind localisation," <https://www.youtube.com/watch?v=VxLjC5co-6s>, 2011.
- [18] P. Sotiropoulos, N. Tosi, F. Andritsos, and F. Geffard, "Optimal docking pose and tactile hook-localisation strategy for auv intervention: The difis deployment case," *Ocean Engineering*, vol. 46, no. 0, pp. 33 – 45, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0029801812000807>
- [19] J. Pineau and S. Thrun, "An integrated approach to hierarchy and abstraction for pomdps," Tech. Rep., 2002.
- [20] L. Mihaylova, T. Lefebvre, H. Bruyninckx, K. Gadeyne, and J. D. Schutter, "Active sensing for robotics - a survey," in *in Proc. 5th Intl Conf. On Numerical Methods and Applications*, 2002, pp. 316–324.
- [21] K. Hsiao, T. Lozano-perez, and L. P. Kaelbling, "Robust belief-based execution of manipulation programs," in *Eighth Intl. Workshop on the Algorithmic Foundations of Robotics*, 2008.
- [22] F. Geffard, P. Garrec, G. Piolain, M.-A. Brudieu, J.-F. Thro, A. Coudray, and E. Lelann, "Tao2000 v2 computer-assisted force feedback telemanipulators used as maintenance and production tools at the areva ncla hague fuel recycling plant," *Journal of Field Robotics*, vol. 29, no. 1, pp. 161–174, 2012. [Online]. Available: <http://dx.doi.org/10.1002/rob.20426>
- [23] Video, "Dof-decoupled active force sensing - box localisation demo," <http://www.youtube.com/watch?v=kcf9SdgqU4I&feature=youtu.be>, 2013.