Exploratory reach-to-grasp trajectories for uncertain object poses

Claudio Zito, Rustam Stolkin, Marek Kopicki, Massimiliano Di Luca and Jeremy L. Wyatt

***Abstract*— This work addresses the problem of planning the reach-to-grasp trajectory for a robotic arm and hand, when there is uncertainty in the pose of the object being grasped. If the object is not in its expected location, then the robot may still gain additional information about the object pose by making tactile or haptic observations if a finger or other part of the hand collides with part of the object during the reach-to-grasp operation. Therefore, it is desirable to plan the reach-to-grasp trajectory in such a way that it takes into account and exploits knowledge about the size and shape of the pose distribution associated with the target pose uncertainty. Here we propose a reach-to-grasp trajectory planning algorithm which addresses this exploration-action problem by trading off a smoothness constraint against likelihood of making haptic observations.**

I. INTRODUCTION

We are interested to generate reach-to-grasp trajectory plans using a fixed-dimensional sampled representation of the uncertainty in pose of the manipulandum (object to be grasped) as suggested in [2]. We wish to build trajectories in a continuous state and action spaces that maximise information gain during reaching-to-grasp, by maximising the difference between observations that would be expected if the manipulandum was in its most likely hypothesis state and the observations that would be expected in any other state, sampled from the distribution of possible states which represents manipulandum pose uncertainty. This approach generates more robust trajectories in the face of uncertainty, by exploring the uncertain region in order to gain information about manipulandum pose while simultaneously attempting to reach the grasping goal state (maximum likelihood hypothesis of manipulandum pose). We have successfully extended this approach up to planning in 6-dimensional spaces, where a 6-DOF robot manipulator equipped with a 15-DOF five-finger humanoid hand plans trajectories that result in an efficient tactile exploration of the uncertain region during reach-to-grasp actions. Note that our entire system is composed by 21-DOF, however we are not interested in changing the configuration of the hand during the reach-to-grasp planning because intuitively we assume that an open configuration of the hand is the most efficient way for a tactile exploration. Nevertheless our planner includes a collision detection and observational model for the entire system as well as an inverse kinematic planner at the end of the reach-and-grasp trajectory in order to close the fingers.

This work is supported by FP7 ICT project GeRT (248273)

C. Zito, R. Stolkin, M. Kopicki, J. L. Wyatt are members of the Intelligent Robotics Lab, School of Computer Science, University of Birmingham, UK (email: {cxz004,stolkinr,msk,jlw}@cs.bham.ac.uk)

M. Di Luca is a member of Centre for Computational Neuroscience and Cognitive Robotics, School of Psychology , University of Birmingham, UK (email: m.diluca@cs.bham.ac.uk)

II. TRAJECTORY PLANER

A. Observational model

The robot should be able to gain additional information about the pose of the manipulandum (object to be grasped) if it makes tactile observations when parts of the hand make contact. We model such observations as the likelihood of reading a contact on the force or torque sensors of the robot, $y_t = h(x_t, p^i), i \in [1, k]$. Mathematically we can define the function $h(\cdot)$ as an exponential distribution,

$$h(x, p^i) = \begin{cases} \eta \exp(-\lambda |x - p^i|_2) & \text{if } 0 < |x - p^i|_2 \leq d_{max} \\ 0 & \text{otherwise} \end{cases}$$

where d_{max} describes a maximum range in which the likelihood of reading a contact is not zero, η is a normaliser and $|\cdot|_2$ is the Euclidian norm on 3D space. Therefore the expected sequence of observations over a trajectory, $u_{1:T-1}$, is:

$$h_t(x, u_{1:t-1}, p^i) = (h(F_2(x, u_1), p^i)^T, h(F_3(x, u_{1:2}), p^i)^T, \\ \ldots, h(F_t(x, u_{1:t-1}), p^i)^T)^T$$

where is the robot configuration at time t if the system begins at state x and takes action $u_{1:T-1}$.

B. Planning a trajectory to maximize information gain

The implementation of our planner uses a Probabilistic Roadmap (PRM) [1] to plan trajectories and detect collisions. Initially a random graph, G, of robot configuration in obstacle-free space is generated. Given a pair x_{root}, x_{goal} which describe the root and goal state of the trajectory, the PRM algorithm finds a joint space trajectory as follows:

- adds x_{root} to G
- finds a obstacle-free configuration $\hat{x}_{goal}$ which is a reachable goal configuration for the robot
- evaluates each node x in the graph G according to

$$c_1(x, x_{root}, \hat{x}_{goal}) = \alpha d(x, x_{root}) \\ + \beta d(x, \hat{x}_{goal}) + \gamma d_{cfg}(x)$$

- starting from x_{root} finds all its neighbouring nodes within a given threshold and compute the follows cost-to-go function

$$c_2(x, x', x_{root}, \hat{x}_{goal}) = \delta d_{bound}(x, x_{root}) \\ + \beta d(x', \hat{x}_{goal}) + \gamma d_{cfg}(x)$$

- finds a path from x_{root} to $\hat{x}_{goal}$ which minimises $c_2(\cdot)$ using A* algorithm

where $\alpha, \beta, \gamma, \delta \in \mathbb{R}$, $d(\cdot)$ is a distance function in $SE(3)$ which combine rotational and transitional distances,

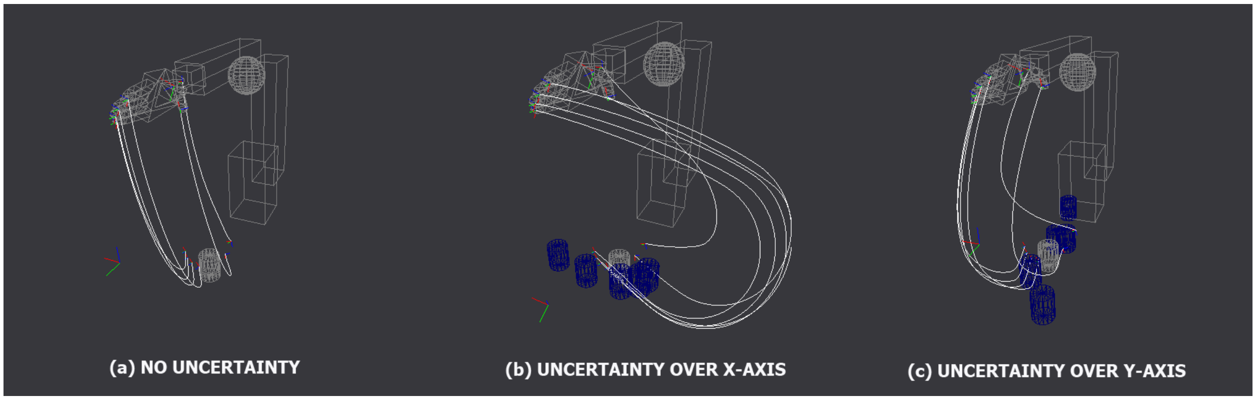

Fig. 1: A 6-axis arm with 15DOF humanoid hand attempts to grasp a cylindrical object under various kinds of uncertainty. White lines denote trajectories of fingertips. White cylinder shows the maximum likelihood pose of the manipulandum. Blue cylinders are particles sampled from the pose distribution of the manipulandum. The image shows three different situations: (a) the system has no uncertainty over the manipulandum pose and plans a smooth, straight trajectory towards the goal state, which is a pre-shaped side grasp; (b) uncertainty prevalently along the x-axis (red axis); (c) uncertainty prevalently along the y-axis (green axis).

$b_{bound}(\cdot)$ is an *ad-hoc* distance function that takes into account also distance in the joint space and has the good property of penalise pair of configurations far away in the joint space and $d_{cfg}(\cdot)$ is a function which penalises dangarous configurations of the robot (i.e. close to joint limits).

Our main contribution is the definition of a new set of heuristics which encode the uncertainty over the object pose. We define,

$$\begin{aligned}\bar{c}_1(x, x_{root}, \hat{x}_{goal}, Q) = &\alpha d(x, x_{root}) \\ &+ \beta d_Q(x, \hat{x}_{goal}) + \gamma d_{cfg}(x)\end{aligned} \tag{1}$$

and

$$\begin{aligned}\bar{c}_2(x, x', x_{root}, &\hat{x}_{goal}, Q, p^{1:k}) = \\ &\delta J(x, u_{1:T-1}, p^{1:k}) d_{bound}(x, x_{root}) \\ &+ \beta d_Q(x', \hat{x}_{goal}) + \gamma d_{cfg}(x)\end{aligned} \tag{2}$$

where Q is the covariance matrix of our sampled states, for any column vector $a \in \mathbb{R}^n$, $d_Q(a) = \sqrt{a^T Q^{-1} a}$ is the Mahalanobis distance, $u_{1:T-1}$ is the trajectory which connect x to x' and $J(x, u_{1:T-1}, p^{1:k}) \in (0, 1]$ is a factor which rewards trajectories with a large difference between expected observations if the object is at the expected location, p^1, versus observations that would be expected if the object is at other poses, $p^{2:k}$, sampled from the distribution of poses associated with the object's positional uncertainty:

$$J(x, u_{1:T-1}, p^{1:k}) = \frac{1}{k} \sum_{i=2}^{k} e^{-\Phi(x, u_{1:T-1}, p^i)} \tag{3}$$

where:

$$\Phi(x, u_{1:T-1}, p^i) = ||h_t(x, u_{1:T-1}, p^i) - h_t(x, u_{1:T-1}, p^1)||^2_{\mathcal{Q}}$$

for each $i \in [2, k]$ and $\mathcal{Q}$ is the block diagonal of measurement noise covariance matrices of appropriate size.

Note that our current observational model is designed to not affect the Eq. 2 when the likelihood of observing a tactile contact is zero. In fact, for robot configurations in which the distance to the sampled poses is larger than a threshold, d_{max}, the cost function $J(\cdot)$ is equal to 1. However we also encode uncertainty in the second factor of our heuristics, $d_{\mathcal{Q}}(\cdot)$, which evaluates the expected distance to the goal configuration. In this way the planner behaves sensibly in the early stages of the trajectory, when the robot is still too far away from the object to observe any contacts.

III. Results

We have tested this algorithm in a simulation environment, in which a 6-axis arm, equipped with a 15 DOF five-fingered humanoid hand attempts to grasp a cylindrical object, subject to various shaped distributions of pose uncertainty. In our experiments we assumed a zero-Gaussian noise with variance of 4 cm along the major axis of noise and 0.25 cm along the minor axis. Fig. 1 shows the resulting reach-to-grasp trajectories produced by our planner. In Fig. 1(a), there is no uncertainty, and so the planner produces the most direct reach-to-grasp path. In Fig. 1(b), there is few uncertainty in the y direction, but lots of uncertainty in the x direction, hence the planner produces a trajectory in which the hand sweeps towards the cylinder along the x axis, thereby maximizing the chance of making contact if the cylinder is not in its expected location. In Fig. 1(c), there is few uncertainty in the x direction but lots of uncertainty in the y direction, hence the hand sweeps towards the cylinder along the y axis, again maximizing the chance of m, should the cylinder not be in its maximum likelihood expected pose.

IV. Acknowledgements

We gratefully acknowledge funding from FP7 ICT project GeRT (248273).

References

[1] L. Kavraki and P. Svestka. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. In *IEEE Transactions on Robotics and Automation*, 1996.

[2] R. Platt, L. Kaelbling, T. Lozano-Perez, and R. Tadrake. Simultanious localization and grasping as a belief space control problem. Technical report, CSAIL, MIT, 2001.