Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Grasp Point Optimization by Online Exploration of Unknown Object Surface

Qiang Li, Robert Haschke, Helge Ritter, Bram Bolder

2012

Preprint:

This is an accepted article published in 12th IEEE-RAS Intl Conf on Humanoid Robots. The final authenticated version is available online at: https://doi.org/[DOI not available]

Grasp Point Optimization by Online Exploration of Unknown Object Surface

Qiang Li, Robert Haschke, Bram Bolder and Helge Ritter
email: {qli,rhaschke,helge}@cor-lab.uni-bielefeld.de bram.bolder@honda-ri.de

***Abstract*—In order to realize in-hand manipulation of unknown objects, we introduce an extension to our previously developed manipulation framework, such that long manipulation sequences, involving finger regrasping, become feasible. To this end, we propose a novel feedback controller, which searches for locally optimal contact points (suitable for regrasping), employing an online exploration process on the unknown object surface. The method autonomously estimates and follows the gradient of a smooth objective function. More concretely, we propose to dynamically switch between manipulability and grasp stability depending on the grasp stability level.**

Physics-based simulation experiments, involving artificial noise to model real-world sensor readings, prove the feasibility of our approach by rotating an object while readjusting the grasp configuration with all fingers in turn.

I. Introduction

A major challenge to exploit the potential of multifingered robot hands for dexterous manipulation is the ability to realize controlled regrasping to enable large-range in-hand object movements. Even if the geometry of the object, the kinematics of the hand, and the friction conditions at all contact points are accurately known, the planning and control of regrasp sequences constitutes a difficult task. Since in most practical cases such information is at best available only partly, recent work has focused on the development of solutions that remain feasible even in the absence of detailed information about the object, the contact friction, and the hand kinematics.

Many of these approaches are based on ideas of robust control, i.e. methods that are little affected by deviations from their underlying model assumptions. Here, we present a novel approach that differs from these lines of research in that it integrates active exploration into the determination of a regrasp sequence: regrasping is split into successive repositionings of a free finger that identifies a suitable next contact point by small exploratory movements across the object's surface in the vicinity of its current contact while simultaneously monitoring a quality measure that combines grasp stability and manipulability in such a way that it can be evaluated under very weak information requirements.

In addition, readily available information from tactile fingertips and joints is used for a feedback controller to compensate deviations between the actual and the modeled motion such that a very simple model with only few assumptions beyond local surface smoothness is sufficient. We demonstrate the feasibility of the approach within a physics-based simulation employing a 22-DOF anthropomorphic Shadow Hand manipulating a spherical object (without using this geometry information in the control method).

The paper is arranged as follows. In section II we discuss how our approach is positioned w.r.t. existing work. In section III, we shortly summarize our manipulation strategy composed of a local manipulation controller and a global finger gait planner as introduced in our previous work [1]. Subsequently, in section IV, we introduce a new control scheme for online exploration of the object surface aiming for an optimization of a given objective function, which is evaluated in section V employing a physics-based simulation. Finally, section VI summarizes our work.

II. State of the Art

Roughly, there are three different lines of research to cope with the dexterous manipulation problem. The first line follows an analytic approach, requiring rather strong assumptions and detailed knowledge about the situation: the hand kinematics, object properties like shape, mass and mass distribution, the contact locations and friction coefficients, and the local surface geometry of both the object and fingertips. Based on this knowledge it is possible to compute joint-level finger trajectories in an offline fashion [2].

The second line is based on the idea of forward control, using a simulation of object-hand interaction [3] to model the grasping and manipulation processes. Grasp poses optimized w.r.t. certain quality criteria [4] become arranged in a pose graph [5] to plan manipulation sequences using state-of-the-art motion planning methods like RRT [6] or PRM [7], tackling e.g. the problem of screwing a light bulb [8].

The third line – in spirit closest to our method – uses feedback as a central mechanism. Ishihara et. al [9] devised a controller that is capable to spin a pen of known shape at an impressive speed. Platt et. al [10] proposed hybrid force/position controllers to realize unknown object grasping by sliding the fingers on the object surface to optimize grasp stability. Tahara et. al [11] point out a method to manipulate objects of unknown shape. They employ a virtual object frame determined by the triangular fingertip configuration of a three-fingered hand to derive a control law to manipulate the object's pose. However, without explicit sensory feedback, their method is limited in accuracy.

III. Object Manipulation Strategy

Reactive feedback-based strategies for object manipulation appear most suitable as a starting point for our approach, because we aim to tackle unknown object manipulation,

Q. Li, R. Haschke and H. Ritter are with the Research Institute for Cognition and Robotics (CoR-Lab), Bielefeld University, Germany. B. Bolder is with the Honda Research Institute Europe (HRI-EU), Offenbach, Germany. Qiang Li gratefully acknowledges the financial support from HRI-EU for the project "Autonomous Exploration of Manual Interaction Space".

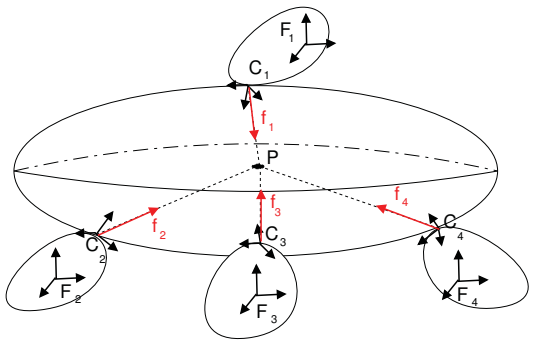

Fig. 1: Force planner employs centroid $\bar{\mathbf{p}}$ of contact locations.

rendering deliberative planning approaches infeasible due to missing object models. Furthermore, we can more directly account for real-life deviations from the planned trajectory: The object pose might be estimated incorrectly, fingers might unpredictably slide or roll or even loose contact at all. To realize manipulation planning in situations where no object model is available, we devise a control strategy to slide an active finger over the object surface to find a new grasp point locally optimizing a given quality measure. Although the gradient-based method cannot find the global optimum, it turns out to be highly sufficient to solve the task.

We divide the manipulation process into two stages: a local manipulation controller and a global finger gait planner. In our previous work we used a set of local controllers from a manipulation control basis [12], [13] to realize small object motions which were sequenced using a finite state machine to realize large-scale in-hand movements of an object [1]. The present paper extends this work with an online control strategy to actively explore the object surface with an active finger, searching an optimal grasp point, while using the other fingers to stably hold the object.

A. Local Manipulation Controller – Position Part

In contrast to traditional planning methods, we aim for unknown object manipulation, expecting as little knowledge as possible. More concretely, we assume point contacts and the availability of coarse contact positions, normals and normal forces, which can be obtained from modern tactile sensor arrays and visual object tracking methods. More detailed properties like friction coefficients are not required.

In order to realize a small object motion $M = O^{-1} \cdot O'$ from the current object pose O to the targeted pose O', we need to determine appropriate finger joint motions. To avoid the need for a detailed object model, we make the essential assumption that contact positions $\mathbf{p}_i^o$ do not move relative to the object within a control cycle. Of course, this is only an approximation. However, the sensor feedback available in the next control cycle will allow us to recognize and correct for undesired contact motion, e.g. due to slipping or rolling. Denoting the current contact positions w.r.t. the object frame by $\mathbf{p}_i^o$, we can easily compute the contact positions $\mathbf{p}_i'$ (w.r.t. the palm) targeted in the current control cycle as follows:

$$\mathbf{p}_i' = O' \cdot \mathbf{p}_i^o = O \cdot M \cdot O^{-1} \mathbf{p}_i \,. \tag{1}$$

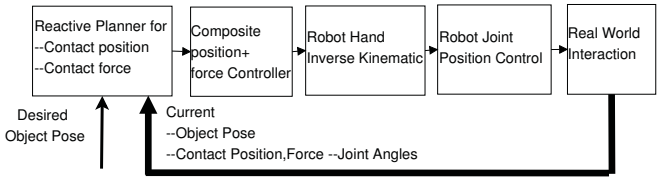

Fig. 2: Low level, local manipulation control scheme.

From this we can compute the required positional changes $\Delta\mathbf{p}_i = \mathbf{p}_i' - \mathbf{p}_i$ for all contact points. Because the object and fingertip geometries are not explicitly taken into account, contacts might have moved due to sliding or rolling after applying the computed hand pose. We might even loose a contact. However, we tolerate such changes and use fast feedback to correct for these unmodeled deviations.

B. Local Manipulation Controller – Force Part

This is achieved by complementing the position control with a suitable force control strategy. Conventional contact force planners strive for a globally optimal contact force distribution to maximize grasp stability under the constraints that (i) all contact forces stay within friction cones, (ii) the total applied force compensates external forces (e.g. gravity), and (iii) local contact forces stay within preset limits.

This general solution is meaningful only if the contact force is controllable. However, we wish to work under the more parsimonious assumption that there is no directional (3D) contact force feedback available, but only a scalar force magnitude along the contact normal.

Following the approach from [11], we plan the force direction such that the resultant moment will be zero by ensuring that the contact force directions of all fingers intersect in one point, which is chosen to be the centroid $\bar{\mathbf{p}}$ of contact locations (cf. Fig. 1). Subsequently we can prescribe force magnitudes along these directions such that the resultant force becomes zero as well. Hence, the force planner calculates desired contact force magnitudes along the *contact directions*, from which we can obtain force errors. Both, the force and positional errors ($\Delta\mathbf{p}_i$) are fed to a composite position/force controller, which calculates the effective contact position error, which in turn is fed to an inverse hand kinematics module to compute joint velocities. To ensure, that the force error converges to zero, we apply a PI-controller for force and a P-controller for position. Fig. 2 summarizes this control scheme. For more details on the local manipulation controller we refer to [12].

C. Regrasp Planner

Joint limits restrict the amount of object motion achievable by a local method. To realize large-scale motions, we need to regrasp before continuing the local manipulation process. Finding a suitable regrasp is the task of the regrasp planner. We use a strategy that is loosely inspired by human manipulation skills: While three passive fingers realize the object motion, the fourth finger takes the role of actively exploring the object to identify an optimal next contact point.

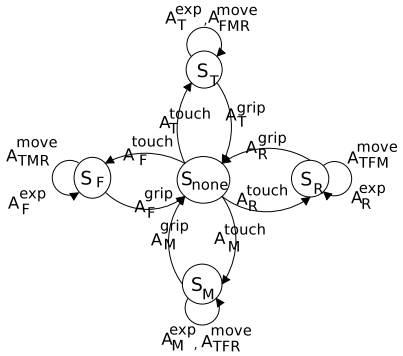

Fig. 3: High level, global regrasp planner.

Subsequently, the fingers switch their active and passive roles and the process is seamlessly repeated. This static finger gait is described by the state machine shown in Fig. 3, ensuring that all fingers explore the object surface, taking the active role in turn. To this end, we distinguish the following states:

S_{none}:	TFMR are all grip fingers	
S_R:	TFM grip fingers	and R exploration finger
S_M:	TFR grip fingers	and M exploration finger
S_F:	TMR grip fingers	and F exploration finger
S_T:	FMR grip fingers	and T exploration finger

Fingers are abbreviated as follows: T: thumb, F: forefinger, M: middle finger; R: ring finger. State transitions are accompanied by a common sequence of object manipulation and exploration motions as follows: After holding the object with all fingers in the initial state S_{none}, an individual finger becomes the actively exploring finger, contacting the object with very small force. In this state, the three remaining, passively holding fingers move the object. After this local manipulation is performed, the active finger slides over the object surface to find the next feasible contact point. Finally, the active finger reestablishes object contact with a contact force determined by the force planner (sec. III-B), thus returning to the state S_{none}. Please note, that the contact force planner considers only the passively grasping fingers and assumes zero contact force for the actively exploring finger. State transitions are triggered by an external process cycling through the states $S_{none} - S_R - S_{none} - S_M - S_{none} - \ldots$

IV. Finding Optimal Regrasp Points

For rotary object manipulation studied in our previous work [1] we could employ precomputed contact points for regrasping. However, for more general object shapes we require a more elaborate search process to find a suitable new contact point to regrasp the object, which subsequently allows to continue the object manipulation motion.

We formalize the contact point selection as an optimization problem by considering two quality criteria, the grasp stability and the manipulability, as an objective function to be maximized in an exploratory search process. Both measures can be easily calculated from the available kinematic and contact information.

A. Grasp Stability and Manipulability

Classical grasp planning theory aims for force-closure grasps which can resist external disturbance wrenches from arbitrary directions without slipping or rolling. Grasp quality measures rely on an analysis of the grasp wrench space, i.e. the set of wrenches applicable to an object through a set of normalized contact forces, which is closely related to the grasp matrix G [2]. The most popular methods employ (i) the minimal singular value of G or (ii) the determinant of G. While the former method yields a worst-case criterion measuring the distance to an unstable grasp configuration along the worst wrench direction, the latter evaluates the volume of the wrench space, which averages over all possible wrench directions. We employ the latter criterion:

$$\phi_{\text{stability}} = \sqrt{\det(G_{\text{passive}} G_{\text{passive}}^t)} \tag{2}$$

Because the exploration phase aims for a recomposition of the group of fingers passively holding the object in the next exploration phase, grasp stability is evaluated considering only contacts of those fingers (which are known beforehand due to the fixed finger gait), which is denoted by G_{passive}.

The manipulability [14] measures the distance of the current hand pose to a singular configuration and thus expresses the capability of the current pose to actively move all fingertips into an arbitrary Cartesian direction. The manipulability is calculated from the Jacobian matrix of the robot hand, which is a block-diagonal matrix formed from the finger Jacobians J_i, assuming uncoupled finger motion. Because during exploration only the Jacobian and thus the manipulability of the active finger changes, we can reduce calculations to the appropriate sub matrix J_i:

$$\phi_{\text{manipulability}} = \sqrt{\det(J_i J_i^t)} \tag{3}$$

Both criteria are complementary: while the grasp stability criterion only considers the contact configuration, the manipulability focuses on the finger motion range, avoiding singular configurations. Hence, both criteria need to be combined effectively. However, a simple linear superposition of both measures is not promising, because often both criteria are conflicting with each other: an increase of stability generates a decrease of manipulability and vice versa, such that a linear superposition would find the least compromise only (see section V for an example). Additionally, it would be difficult to find a suitable weighting of both components, because they are not normalized.

Hence, we employ a hierarchical combination of both criteria, which are chosen in an adaptive fashion. Because we aim for in-hand object manipulation, we propose to normally use manipulability as the primary criterion which is optimized as long as the grasp stability criterion fulfills a given minimal threshold. Hence, if the grasp stability is larger than the threshold, we generate active finger motion following the estimated gradient of the manipulability, thus maximizing this criterion. Otherwise, if grasp stability is below the threshold, finger motion will be generated along

the gradient of grasp stability in order to increase stability to the necessary level first.

B. Object Surface Exploration Controller

Given a selected objective function ϕ, the task of the exploratory motion controller is to generate a sliding motion over the (unknown) object surface which finds and follows the objective's gradient $\nabla\phi$. Because both quality criteria are complex non-linear functions of the contact point motion $\dot{\mathbf{c}}$, we do not try to find a closed-form solution of the gradient, but aim for its online estimation.

For a number n of control cycles, the exploratory motion follows the estimated gradient $\tilde{\nabla}\phi$, which is projected onto the tangential plane at the current contact point:

$$\dot{\mathbf{c}}^r = \eta \cdot T_c^r \cdot P(\mathbf{c}) \cdot (T_c^r)^{-1} \cdot \tilde{\nabla}^r \phi\,, \tag{4}$$

where η is the gain factor, T_c^r the homogeneous transformation from the contact frame (subscript c) to the global reference frame (subscript r), and $P(\mathbf{c}) = \mathrm{diag}(1,1,0,1)$ a projector mapping onto the tangent plane at the contact point $\mathbf{c}$. This targeted contact velocity is fed into a hybrid position/force controller (cf. [12]), which tries to realize this motion in the tangent plane while simultaneously maintaining a (small) normal contact force onto the object. After n control cycles we re-estimate the gradient $\nabla_{\mathbf{c}}\phi$ from the slope of the objective function during motion $\Delta\mathbf{c}$ from $\mathbf{c}$ to $\mathbf{c}'$. This procedure is iterated until convergence or until a predefined number of $N_{\max}$ control cycles have been executed.

The complete control algorithm is summarized in Alg. 1. In order to obtain a smooth estimation of the gradient, in line 8 we apply a sliding average using a smoothing coefficient of λ=0.9. Please note, that the estimated gradient $\tilde{\nabla}\phi$ is represented w.r.t. the time-constant global reference frame. The contact reference frames which are used to compute $\nabla\phi$ change over time and thus are not suitable for a time-consistent representation.

Further note, that the targeted motion is restricted to the tangent plane of the contact point in order to maintain contact. Of course, due to approximation errors and foremost due to the unknown object shape, the real contact motion might also have a normal component, thus changing the contact force or loosing contact at all. However, the hybrid position/force controller accounts for these deviations from the planned motion by maintaining a given contact force. Obviously this control algorithm only assumes that the object surface changes smoothly along every contact trajectory, such that gradient estimates can be computed and small motion deviations along the surface normal can be compensated.

Algorithm 1 Object surface exploration, maximizing ϕ

1: $i = 0$ {initialize cycle count}
2: $\tilde{\nabla}\phi^r \propto \mathcal{N}(0, \sigma = 0.15)$ {randomly initialize gradient}
3: **while** $++i \leq N_{\max}$ **and** $\|P \cdot (T_c^r)^{-1} \cdot \tilde{\nabla}\phi^r\| \not\leq \varepsilon$ **do**
4: **if** $i \mod n = 0$ **then** {update $\tilde{\nabla}\phi^r$ every n cycles}
5: $\Delta\mathbf{c}_c = P(\mathbf{c}) \cdot T_r^c \cdot (\mathbf{c}'_r - \mathbf{c}_r)$
6: $\nabla\phi^c = [\frac{\phi'-\phi}{\Delta\mathbf{c}_x^c}, \frac{\phi'-\phi}{\Delta\mathbf{c}_y^c}, 0]^t$
7: limit norm of $\nabla\phi^c$ to $\nabla_{\max}$
8: $\tilde{\nabla}\phi^r \leftarrow \lambda \cdot \tilde{\nabla}\phi^r + (1-\lambda) \cdot T_c^r \cdot \nabla\phi^c$
9: $\phi' \leftarrow \phi\,, \quad \mathbf{c}' \leftarrow \mathbf{c}$
10: **end if**
11: $\dot{\mathbf{c}}^r \leftarrow \eta \cdot T_c^r \cdot P(\mathbf{c}) \cdot (T_c^r)^{-1} \cdot \tilde{\nabla}\phi^r$
12: **end while**

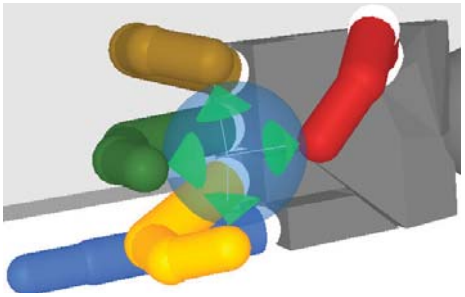

Fig. 4: Shadow Hand model rotating a sphere using 4 fingers.

V. Simulation and Discussion

The object manipulation algorithm is validated in a physical simulation experiment. We use the Vortex physics engine to obtain real-time contact information (contact position and normal force magnitude) and the object's pose (position and orientation). The former information will also be accessible in real world, exploiting modern tactile fingertip sensors providing a moderate spatial resolution. Exploiting the known sensor shape and kinematic model of the hand, we can calculate contact positions relative to the hand and correlate them to a visually tracked, coarse object model.

In order to model noisy real-world sensors, artificial white noise is superimposed on the feedback provided by the physics engine. Particularly, the standard deviations of added Gaussian noise for the contact positions is 0.3cm, and 0.1 for the contact forces. Hence, the positional noise resembles the spatial resolution of tactile sensors in the Shadow Hand. Note that we do not model calibration errors, which usually result in systematic deviations.

As the algorithm assumes a smooth object surface, we test one exemplary object – a sphere of 5cm diameter, which has to be rotated in-place by a 22-DoF Shadow Hand model (see Fig. 4). Note, that the shape of the object as well as other parameters like friction properties are not available to the manipulation strategy. For initial grasping, we assume that the object is located in a suitable pose relative to the hand and desired grasp points are known. The simulation is resembling our real robot setup to facilitate future transfer into real world, once the required tactile feedback is robustly available from fingertip sensors. The whole manipulation process comprises three phases:

- Grasp the object while it is fixed in the world, which is necessary to avoid that the object is kicked off.
- Unfreeze the object and stabilize the grasp employing active force control in order to prepare manipulation.

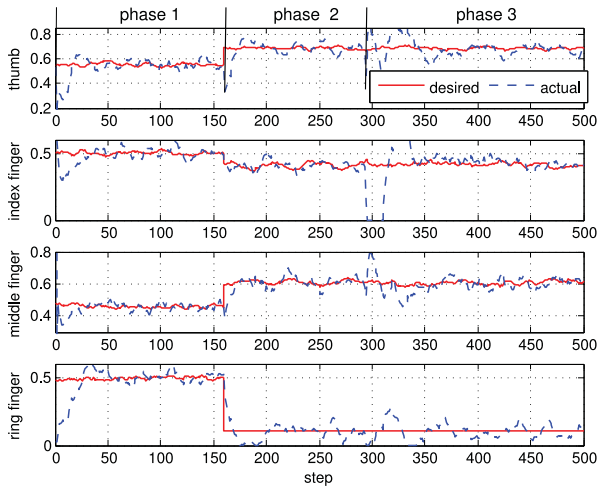

(a) Contact force magnitudes along the contact normals

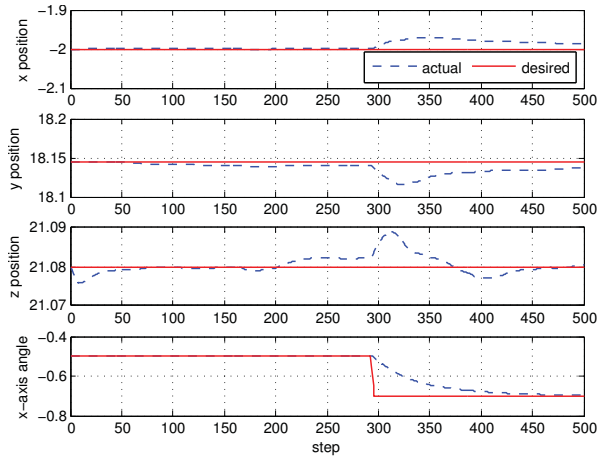

(b) Object pose: x, y, z position and x-axis rotation

Fig. 5: Force and motion trajectories while rotating the object

- Actually manipulate the object:
  - Choose an active finger, releasing the grip, i.e. reducing contact force (transition to state S_X).
  - Rotate the object by a small amount, e.g. $5°$ using the three remaining, passive fingers.
  - Explore the object surface according to algorithm 1.
  - Reestablish grip with all fingers (trans. to S_{none}) and continue manipulation with next active finger.

We focus the evaluation on the last item, the active exploration process to find improved contact points. The realization of a stable grasp, which is the objective of the first two stages, is reported in our previous work [12], [15].

a) Relaxing grip and rotating the object: The first experiment shown in Fig. 5, illustrates the actions A_R^{touch} and A_{TFM}^{move}, i.e. the ring finger relaxing the grip (phase 2), and subsequent rotation of the object around the x-axis using the remaining fingers (phase 3). The first phase shown in the figure corresponds to the state S_{none}, where all four fingers hold the object employing the active force controller described in sec. III-B.

In Fig. 5a, the contact force evolution during all three phases is shown, while Fig. 5b shows the motion trajectories of the object position (x, y, z) and its rotation angle around the x-axis. Target trajectories are visualized as red solid lines, while actual trajectories are depicted as blue dotted lines. Note that the noisy force trajectories are mainly due to the artificially superimposed sensor noise, which should model real-world conditions.

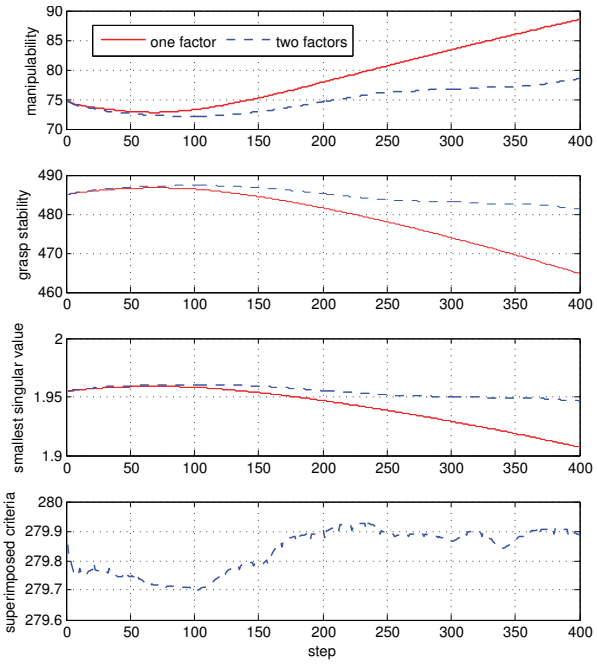

Fig. 6: Evolution of quality measures when ring finger is exploring the object surface: manipulability, average grasp stability, smallest singular value of G, superimposed criteria.

At the transition from phase 1 to phase 2, the desired contact force for the ring finger (last sub graph) is lowered to 0.1. The other fingers slightly adapt their contact forces to account for the omission of the fourth contacting finger. The transition to phase 3 doesn't change the force profile, but only realizes the object rotation (cf. [12]). The simulation results show, that the object position error is less than $0.1cm$, and the orientation error is less than $0.01rad$.

b) Object surface exploration: Fig. 6 shows the evolution of the quality measures during two different exploratory motions of the ring finger (action A_R^{exp}): The red solid lines result from the motion, when only the manipulability is optimized, while the blue dotted lines result from the motion when both quality criteria are optimized simultaneously using a linear superposition with weights $\frac{1}{2}$ each. Both motions start from the same initial configuration.

As can be seen from the graphs, both criteria are conflicting with each other, resulting in only minor changes to the overall quality measure in the superimposed case. However, when the optimization focuses onto a single criterion, the proposed control strategy can successfully estimate the gradient direction after a few iteration steps and subsequently follow this gradient for maximization. The decrease of the manipulability during the first 75 time steps is due to a

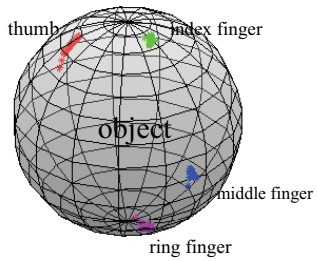

(a) Contact point cloud generated by object exploration

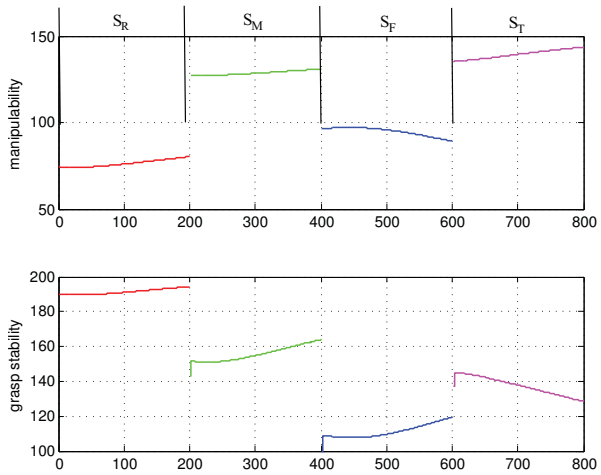

(b) Evolution of manipulability and grasp stability

Fig. 7: Results of a complete manipulation sequence involving all states, i.e. exploring with all fingers in turn.

poor random initialization of the gradient direction (line 2 of alg. 1), which is only slowly overcome due to the sliding average (line 8). As can be seen from the stability graphs, the grasp is stable during the course of manipulation, although stability may decrease. Repeating the exploratory motion 50 times using different initial gradient directions, always leads to a successful maximization of the objective function.

c) Continuous Object Manipulation: Finally, Fig. 7 shows the evolution of the contact points w.r.t. the object frame (a) and the corresponding quality criteria (b) during a continuous rotation of the object following the finger gait pattern as defined by the state machine in Fig. 3: exploring with ring, middle, index finger, and thumb in this order.

As can be seen from the evolution of the objective functions, the algorithm always maximizes the manipulability of the exploring finger – except in state S_F, where the grasp stability is chosen as the objective function. In this state, the grasp stability – evaluated for the group of passive fingers – drops below the chosen threshold, because the thumb becomes active and is thus excluded from the holding task. The complete manipulation sequence is illustrated in the accompanying video to this paper.

Please notice, that the exploration process also reveals valuable shape information of the object as illustrated by the contact point cloud in Fig. 7a. If object manipulation is continued, more and more contact points are sampled and can serve as a basis to reconstruct the object shape [16].

VI. Summary and Outlook

The proposed, novel control algorithm to search the maximum of a given smooth objective function in an exploratory motion process, sliding a fingertip over an unknown object surface, provides another missing puzzle piece to realize complex in-hand manipulation.

Our algorithm makes as little use of prior knowledge as possible. Neither the object shape, nor detailed contact properties need to be known. The controller relies on tactile feedback to obtain the current contact point, contact normal (which is known from the local finger geometry), and contact force magnitude. Vision feedback can be employed to track the object pose. Due to its closed-loop characteristics, the approach is very robust to sensor noise and small deviations from the planned motion, which was demonstrated in physics-based simulation experiments.

In future, we will speed up the gradient estimation process, which is slowed down using averaging methods to overcome random fluctuations of the objective functions. Also, we will apply the algorithm to more complex object shapes, which is easily possible, because only smoothness constraints have to be fulfilled.

References

[1] Li, Haschke, Ritter, and Bolder, "Rotary surface object manipulation by multifingered robot hand," in *7th German Conf. for Robotics*, 2012.
[2] R. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.
[3] Miller and Allen, "Graspit! a versatile simulator for robotic grasping," *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, 2004.
[4] A. Miller and P. Allen, "Examples of 3D grasp quality computations," in *Proc. ICRA*, vol. 2, 1999, pp. 1240–1246.
[5] T. Phoka and A. Sudsang, "Contact point clustering approach for 5-fingered regrasp planning," in *Proc. IROS*, 2009, pp. 4174–4179.
[6] M. Yashima, "Manipulation planning for object re-orientation based on randomized techniques," in *Proc. ICRA*, vol. 2, 2004, pp. 1245–1251.
[7] J. Saut, A. Sahbani, S. El-Khoury, and V. Perdereau, "Dexterous manipulation planning using probabilistic roadmaps in continuous grasp subspaces," in *Proc. IROS*, 2007, pp. 2907–2912.
[8] Z. Xue, J. Zollner, and R. Dillmann, "Dexterous manipulation planning of objects with surface of revolution," in *Intelligent Robots and Systems, IEEE International Conference on*, 2008, pp. 2703–2708.
[9] T. Ishihara, A. Namiki, M. Ishikawa, and M. Shimojo, "Dynamic pen spinning using a high-speed multifingered hand with high-speed tactile sensor," in *Proc. Humanoids*, 2006, pp. 258–263.
[10] R. Platt Jr, "Learning and generalizing control-based grasping and manipulation skills," Ph.D. dissertation, UMass Amherst, 2006.
[11] K. Tahara, S. Arimoto, and M. Yoshida, "Dynamic object manipulation using a virtual frame by a triple soft-fingered robotic hand," in *Proc. ICRA*, 2010, pp. 4322–4327.
[12] Q. Li, R. Haschke, H. Ritter, and B. Bolder, "Simulation results for manipulation of unknown objects in hand," in *Int. Conf. on Robotics and Biomimetics*, 2011.
[13] Q. Li, M. Meier, R. Haschke, H. Ritter, and B. Bolder, "Object dexterous manipulation in hand based on finite state machine," in *IEEE Int. Conf. on Mechatronics and Automation*, 2012.
[14] Y. Nakamura, *Advanced Robotics: Redundancy and Optimization*. Addison-Wesley, 1990.
[15] Röthling, Haschke, Steil, and Ritter, "Platform portable anthropomorphic grasping with the bielefeld 20-dof shadow and 9-dof tum hand," in *Proc. IROS*, 2007, pp. 2951–2956.
[16] M. Meier, M. Schöpfer, R. Haschke, and H. Ritter, "A probabilistic approach to tactile shape reconstruction," *IEEE Transactions on Robotics*, vol. 27, no. 99, pp. 1–6, 2011.