

Bio-inspired Motion Strategies for a Bimanual Manipulation Task

Jan Steffen, Christof Elbrechter, Robert Haschke and Helge Ritter
Neuroinformatics Group, Faculty of Technology, Bielefeld University, Germany
{jsteffen, celbrech, rhaschke, helge}@techfak.uni-bielefeld.de

Abstract— We consider the complex task of coordinating two five-fingered anthropomorphic robot hands for taking a jar passed from a human user and unscrewing its cap. Using a pair of 7-DOF redundant arms for operating the hands, we study how the incorporation of human movement strategies at the finger and arm levels can aid in the solution of the overall bimanual task. At the finger level, we employ a finger control manifold for the unscrewing motion that has been synthesized with a kernel approach applied to human motion data captured with a data glove. At the arm level, we use a heuristic motivated from the observation of human arm movements to enhance the space of pass-over configurations that the system can successfully handle. In addition, we provide a brief description of the architecture of the overall system that comprises 54 motor degrees of freedom and integrates camera vision, arm and finger control as well as a speech output component for interaction with the human user.

I. INTRODUCTION

The coordination of multifingered robot hands for accomplishing "ordinary" human manual actions belongs to the most demanding challenges in robotics. Contributing factors include the large number of involved motor degrees of freedom, the need to cope with redundancy, variable contact geometries in the face of imprecise object geometry, uncertainty in the actuator characteristics, the lack of adequate tactile feedback, the need to integrate vision, and many more.

Yet, surmounting this challenge is of paramount importance if robots are to become useful in many common household settings where settings are tailored to the characteristics and capabilities of human hands. Besides these application considerations, given the impressive range of human manual actions, their replication with anthropomorphic robot hands is likely to be both a strong driving force and a severe "litmus test" for a thorough understanding of a significant part of our sensory-motor cognition.

This background is the rationale for a major long-term research effort at the lab of the authors, where the exploration of bio-inspired motion strategies and their integration in order to realize a manipulation system that can cope with a sizable range of manual interaction tasks is a primary focus.

The present paper focuses on the combination of such motion strategies to coordinate arm and finger motions for a major generic skill: namely, the opening of containers. In humans, this skill is a complex capability developed over years of sensory-motor learning. For a robot, it represents a major next step beyond grasping, which itself constitutes one element of container opening and for which a considerable number of quite successful approaches have already been developed and published in the literature [1]–[10].

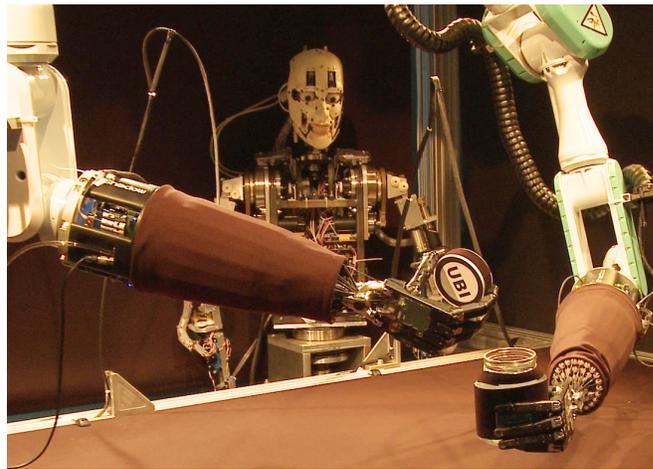


Fig. 1. Considered scenario for bimanual manipulation: The cap of a glass jar received from human should be unscrewed.

Compared to grasping, opening of containers requires the robot to master the coordination of a larger number of subtasks, including the recognition of the object, the discrimination of its essential components, the type of required grasping and opening actions, and the acquisition, representation and generation of the manipulation movements themselves. Moreover, while grasping of typical objects is usually feasible with a single hand, opening a container in most cases requires bimanual activity, with one hand fixating the object and the second performing the opening action.

In this paper, we demonstrate how a combination of three bio-inspired approaches can cover a significant portion of the involved coordination tasks, such that the flexibility of the system to cope with object configurations is increased and part of the burden of prestructuring manipulation motions is replaced by learning from human movement data. This approach follows the spirit of earlier approaches that used bio-inspired methods and learning strategies, such as imitation learning and learning from demonstration [11], [12], to provide the system with increased flexibility. So far, the majority of these approaches have focused on "large scale" movements at the body or entire arm levels. The present work extends this focus into the realm of fine manipulation skills which so far has received comparably less attention along the lines of bio-inspired algorithms.

Specifically, we propose a strategy motivated by the observation of human arm movements and exploiting the rotational symmetry of a to-be-grasped object for choosing a favorable configuration of a 7-DOF redundant robot arm. We

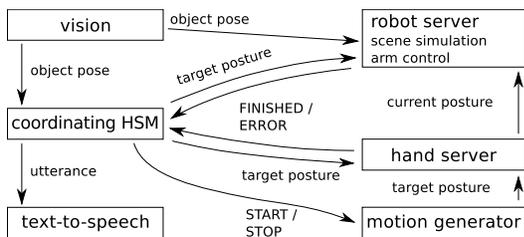


Fig. 2. System overview: distributed processes for vision, robot control and speech output are coordinated by a central hierarchical state machine (HSM) dispatching system events.

demonstrate that our approach significantly enlarges the set of feasible object poses that can be handled by the system, when compared to a standard method.

Coordination at the arm level is complemented by coordination at the finger level, utilizing a kernel-based manipulation manifold extracted by a learning algorithm from human movement data captured with a data glove. Our previous results on this approach have been restricted to the context of simulation settings [13]. The present contribution is the first presentation of this method for a complex multifingered manipulation movement synthesized from human training data and performed on a bimanual arm-hand system using a pair of anthropomorphic robot hands.

These two approaches and their combination are described in Secs. IV-VI and constitute the main new contribution of this paper. The third bio-inspired motion concerns the grasping phase, where we follow a tactile-driven approach developed in previous work [6]. To make the paper more self-contained, it is briefly sketched in Sec. III.

As a benchmark example of realistic complexity, we focus on the case of opening a glass jar that is known to the robot. A vision system provides the necessary position information when a human passes the jar to the robot, after which the robot has to autonomously perform all the necessary actions to unscrew the lid. A dialog component [14] gives verbal feedback when the object is not properly visible or is relocated during the grasping phase.

A similar task has already been successfully demonstrated by the DLR group using the humanoid platform Justin: they successfully demonstrated the opening of a screw cap bottle [15]. During the manipulation, which is preprogrammed for their four-fingered robotic hand, the cap is perfectly grasped and unscrewed. While this is an impressive feat, we demonstrate that a similar capability can also be realized by a combination of several biologically motivated movement primitives. Moreover, an essential part of the action – the fine manipulation of the jar lid – is synthesized from a generalizable learning algorithm in our approach.

An inevitable overarching concern has been the integration of all required functional modules. This poses an additional challenge in itself and is beyond the scope of the present paper. However, to give the reader a feeling of how the above movement primitives have been implemented in the context of the overall system, we briefly describe the underlying integration architecture in Sec. II (for details, see [16], [17]).



Fig. 3. For object detection and localization a topological marker is used, which easily allows to extract an object reference frame. Left-most image shows result of local threshold operator.

II. SYSTEM OVERVIEW

The Bielefeld “Curious Robot” Setup [18] combines visual attention and an intelligent interface for speech recognition and understanding with dexterous bimanual manipulation skills and proactive dialog communication in order to study multi-modal task-oriented instructions given by lay persons. For manipulation tasks, the setup employs two redundant 7-DOF Mitsubishi PA-10 robot arms as well as two 20-DOF Shadow Dexterous Hands, accumulating to 54 DOF within the whole system. The robot arms are mounted from the ceiling providing a humanoid arm configuration. They are controlled by a robot server realizing a security concept for collision detection based on an internal scene model.

The Shadow Dexterous Hands [19] are distinguished by their human-like design: in size, number and flexibility of joints the hands resemble their human counterpart in a unique manner. Their actuation with McKibben style pneumatic muscles on the one hand dramatically eases grasping tasks due to the inherent compliance of the fingers. On the other hand, complex nonlinear effects like hysteresis, friction and compressibility of the air render the generation of smooth finger motions a difficult task.

The whole robot system is controlled by numerous distributed processes tied together using the XCF middleware toolkit [17] which features an event-driven communication scheme. The coarse structure of the system is shown in Fig. 2. A central scene simulation component, located in the robot server process, is used for collision detection based on a prediction of future motion of all robot components. To maintain an up-to-date scene description, the hand server publishes its current hand posture. The vision module is responsible for recognizing and localizing the object. To simplify the vision task, we employ a visual marker on the object as shown in Fig. 3. Finally, the time flow of actions is coordinated by a central hierarchical state machine (HSM), which monitors the current system state and chooses an appropriate action in response to an incoming system event.

III. PROTOTYPE BASED GRASP FORMATION

Dexterous robotic grasping has been an active research subject for decades, and many grasp planning algorithms have been proposed. Finding feasible grasp postures requires a search in the high-dimensional space composed of the internal degrees of freedom of the robotic hand as well as its position and orientation relative to the object. Many grasp planning approaches employ an offline optimization process to generate a database of feasible grasp postures,

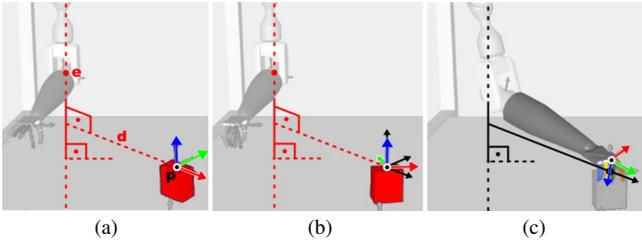


Fig. 4. Determining an arm posture exploiting the symmetry of the cap. (a) The relative object-hand transform associated to a grasp type is recorded for a normalized object pose, where the x -axis (red) is aligned to the vector \mathbf{d} pointing from the elbow to the object. (b) For a new object pose a normalized frame (black) is computed according to Eq. 1. The object’s normal axis (blue) is maintained. (c) The realized arm pose reduces the motion of the elbow joint.

which are subsequently used to look up the best-matching grasp for a specific real-world object. The widely used grasping simulator *GraspIt!* [20] generates candidate grasps by continuously closing all fingers starting from a given pre-shape posture until object contact is established. This procedure is repeated for a set of approach directions, which are regularly spread around the object [2] or are estimated from 3D point clouds describing the object geometry [3].

While this common grasping approach generates nearly optimal grasp postures, it has a major drawback: it assumes that the shape of the object to be grasped is available for the optimization process and can be matched to an object shape in the database. In contrast to this deliberative approach are reactive grasping strategies exploiting tactile feedback and compliance to realize stable grasps for even unknown, or only coarsely known object shapes [4], [6]. Our grasping strategy falls into the latter category. Based on the coarse shape of the object, i.e. its size and elongation, which can be easily determined from vision, we choose a grasp prototype from a set of three basic prehensile grasps suggested from grasp taxonomies like [21]. In particular we distinguish power, pinch, and precision grasps. Each prototype comprises a pre-shape and a target hand posture as well as the position and orientation of the hand relative to the object. The grasp is realized by an enclosing finger motion from the pre-shape towards the target posture until the fingers exert sufficient force on the object. Due to the size and large weight of the jar, we only employ a power grasp in the scenario at hand.

IV. SYMMETRY-BASED REDUNDANCY RESOLUTION FOR ENHANCING THE SET OF FEASIBLE OBJECT POSES

Both described grasping strategies require to specify the approach direction of the grasp given as the position and orientation \mathbf{T} of the hand relative to the object frame. The first strategy finds a feasible end-effector pose from a thorough search within the database, while the latter strategy utilizes a fixed transformation relative to the object frame which is extracted from vision in both cases. For the bio-inspired method, this determines a *sole* hand pose to realize the grasp, thus dramatically limiting the set of object poses which can be successfully grasped. However, many everyday objects exhibit symmetries which can be exploited in order to choose the approach direction in a more flexible manner,

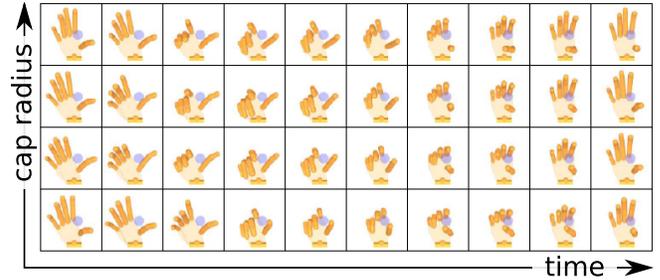


Fig. 5. Visualization of a structured UKR latent space for the cap turning movement. The hand pictures correspond to the mappings of a regular grid in latent space back into the observation space of hand postures. The horizontal latent dimension corresponds to the temporal component of the represented movement and the vertical latent dimension controls the motion parameter (the cap radius). As comparison aid, a small blue cap of radius $r = 1.5\text{cm}$ is depicted in each sub-picture.

eventually optimizing subsidiary goals, e.g. maintaining a natural pose, avoiding joint limits or obstacles.

From observation of human reaching movements, we derived a simple heuristic to resolve the ambiguity of the free rotational DOF around the symmetry axis of the cylindrical jar top: The forearm should be aligned relative to the object such, that the deviation of the elbow from its *initial* position \mathbf{e} in relaxed posture is minimized (see Fig. 4). Given the original object frame $\mathbf{O} = [\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{p}]$ obtained from vision we compute an adapted object frame $\mathbf{O}' = [\mathbf{y}' \times \mathbf{z}, \mathbf{y}', \mathbf{z}, \mathbf{p}]$ which corresponds to a normalized orientation of the object using a *virtual* rotation about its symmetry axis \mathbf{z} , i.e. the normal axis of the jar cap.

In normalized pose, the projection of the \mathbf{x}' -axis to the table surface ($z = 0$) is parallel aligned to the vector \mathbf{d} pointing from the elbow origin \mathbf{e} to the cap center \mathbf{p} , again projected to the table surface. This is accomplished with the following equations:

$$\mathbf{y}' = \frac{\mathbf{z} \times \mathbf{d}}{\|\mathbf{z} \times \mathbf{d}\|} \quad \text{where} \quad \mathbf{d} = (\mathbf{p} - \mathbf{e})|_{z=0} \quad (1)$$

Applying the fixed object-hand transformation \mathbf{T} associated to a particular grasp type we obtain the final tool transform $\mathbf{T}_{tool} = \mathbf{O}' \cdot \mathbf{T}$ for grasping.

The proposed heuristic to find an end-effector pose suitable for grasping is not only useful for our bio-inspired reactive grasping strategy. Rather, it can also be used to bias the search for a feasible grasp in the database approach, which may significantly boost the online performance of those methods.

V. EXTRACTING A MANIPULATION MANIFOLD FROM HUMAN TRAINING DATA

To represent and learn complex motion trajectories from observation, we previously proposed *Structured UKR manifolds* [22], [23] which are based on *Unsupervised Kernel Regression* (UKR [24]), a recent approach to learn non-linear continuous manifolds. UKR finds a low-dimensional (latent) representation $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \subset \mathbb{R}^{q \times N}$ of observations $\mathbf{Y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N) \subset \mathbb{R}^{d \times N}$ together with a functional relationship $\mathbf{y} = \mathbf{f}(\mathbf{x})$. In particular, the method

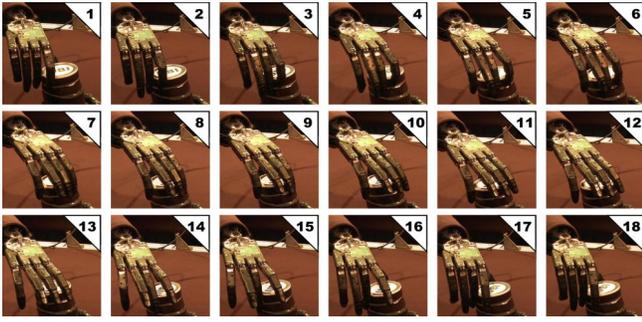


Fig. 6. One cycle of the cap turning manipulation.

uses an unsupervised generalization of the Nadaraya-Watson kernel estimator

$$\mathbf{f}(\mathbf{x}; \mathbf{X}, \mathbf{Y}) = \sum_{i=1}^N \mathbf{y}_i \frac{K(\mathbf{x} - \mathbf{x}_i)}{\sum_j K(\mathbf{x} - \mathbf{x}_j)}. \quad (2)$$

to find optimal latent variables \mathbf{x}_i which minimize the reconstruction error

$$R(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{y}_i - \mathbf{f}(\mathbf{x}_i; \mathbf{X}, \mathbf{Y})\|^2. \quad (3)$$

$K(\cdot)$ represents a multivariate (Gaussian) kernel with bandwidths $H = 1$. In our context, observations \mathbf{y}_i are sequences of observed human hand postures realizing a certain task. A suitable latent space \mathbf{X} for trajectories is time. Other latent dimensions can reflect additional task parameters. In the considered screwing task we use the cap radius as an parameter which defines the average opening of the hand.

Structured UKR [23] extends the original cost function (3) with an additional term to penalize a disordering of data sequence elements along the time dimension. To learn periodic motions, a periodic kernel $K(\cdot)$ is used. It was shown, that *Structured* UKR is capable to learn smooth manipulation manifolds from observed training trajectories in a robust manner.

We obtained suitable training data by recording sequences of hand postures (in the form of joint angle vectors) from demonstrations with a hand-mounted dataglove. Due to the mismatch of the kinematic structure of the human and the robotic hand, the angular values measured by the dataglove cannot directly be applied to control the robot hand. Rather, the demonstrator used the glove as an input device for a *hand model* executing the screwing task *in simulation*. Employing visual feedback the human is able to learn an adequate hand posture mapping to successfully achieve the task. Finally, the joint angles of the simulated hand, which incorporate contact feedback from an underlying physics-engine, can be used to control the real robot hand.

Using this indirect method that follows the ‘Robot Skill Synthesis via Human Learning’ paradigm [25], we recorded sequences of hand postures during cap turning movements for five different cap radii ($r = 1.5\text{cm}$ to 3.5cm). For each of these radii, we produced five sequences each of about 30 to 45 hand postures. Each hand posture consists of a vector of the 24 joint angles of the simulated hand model.

The UKR manifold of hand postures is spanned by a two-dimensional latent space as shown in Fig. 5: the first (periodic) latent dimension encodes the temporal component of the motion and the second one the corresponding cap radius. The manifold smoothly generalizes also for cap radii not seen during training, which motivates the employment of the rather costly UKR approach to manifold learning.

Navigating through the latent space and applying the manifold mapping $\mathbf{f}(\cdot)$ generates smooth trajectories in hand-posture space. Maintaining a specific latent radius and progressing along the latent time dimension results in hand posture sequences, which resemble the trained turning motion for the given cap radius. Whereas this control scheme has been presented already in [22], [23], this is the first time that it has been applied to a real robot hand.

In preparation of the manipulation, a predefined relative object-hand position is actuated, again exploiting the symmetry of the jar cap to find an optimal end-effector pose as described in the previous section. Subsequently, the latent trajectory starts from an empirically determined initial position \mathbf{x}_0^* in latent space, which maps to an open hand posture $\mathbf{y}_0^* = \mathbf{f}(\mathbf{x}_0^*; \mathbf{X}, \mathbf{Y})$. To actually maintain contact to the jar cap during the turning motion, we have chosen the latent radius to be slightly smaller than the real cap radius. Employing the passive compliance of the Shadow Hand this guarantees firm contact without requiring a complex tactile feedback controller.

By increasing the temporal latent variable and maintaining the chosen radius, a consecutive series of latent target parameters ($\mathbf{x}_0^*, \mathbf{x}_1^*, \mathbf{x}_2^*, \dots$) with corresponding mappings ($\mathbf{y}_0^*, \mathbf{y}_1^*, \mathbf{y}_2^*, \dots$) is generated and online sent to the hand server to actuate the robot hand. Each \mathbf{y}_i^* constitutes an intermediate hand posture of the represented turning motion and the consecutive actuation of the ongoing series of \mathbf{y}_i^* s therefore realizes the targeted motion.

Figure 6 depicts a series of key frames of the resulting motion performed by the real Shadow Hand as part of the presented bimanual task (see accompanying video [26]). The motion consists of three phases: Frames 1 to 6 show the turning phase, in which thumb, index finger and middle finger have contact with the cap and the cap is being rotated. Frames 7 to 12 show the opening phase, in which the fingers move away from the cap in order to release the finger contacts. Frames 13 to 18 show the backward movement of the hand with opened fingers until the initial hand posture of Frame 1 is reached and the cycle restarts.

VI. CASE STUDY: CAP TURNING MOVEMENT

To develop the previous elements within a concrete benchmark scenario, we considered the case of the cap turning movement for the opening of a glass jar. This type of movement represents a major representative of the larger class of container opening tasks. Variants of it occur in many daily tasks; moreover, its (approximate) periodicity is another generic feature of many interesting types of manipulation movements.

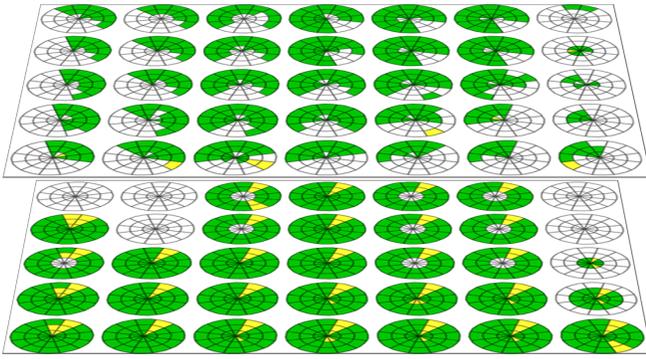


Fig. 7. Visualization of feasible object poses for pick-up using the classical (top) and optimized (bottom) computation of tool frame. The centers of the discs are located at the evaluated object positions, the segments indicate its orientation. Each disc is divided into three rings, corresponding to the object elevation above the table ($z=0/15/30cm$ for the outer, middle and inner ring resp.). Green color indicates, that both picking the object and hand-over have been successful. White indicates, that picking already failed, while yellow color is used, when picking succeeded, but subsequent hand-over failed, because the object could not be moved to the fixed hand-over pose.

For the experiments, we used a medium size nutella[®] jar with a plastic lid whose surface was chosen to offer good contact friction with the finger tips. When the jar was passed to the robot, it had to grasp the rotationally symmetric lid with its right hand, employing the hand positioning strategy described in Sec. IV. Grasping the lid allowed the left hand to grasp the jar body, thereby freeing the right hand. This hand-over was realized at a favorable (and preprogrammed) standard position. Notice, that the jar body is *not symmetric* but rather oval (as can be seen from Fig. 3), such that the full 3D orientation of the jar has to be considered. The unscrewing operation was then initiated by the now free right hand starting from an end-effector position computed from an updated jar localization in order to correct for errors during the hand-over process.

The unscrewing operation consisted of the repetitive cycling through the motion manifold trained from human demonstration data as described in the previous section. The robot was able to unscrew the jar in the large majority of cases (18 out of 20). A small fraction of trials failed due to blocking of the jar, or because an accumulation of undesired finger movements in the supporting left hand brought the jar into a position in which a continuation of screwing was no longer possible.

Even in the large majority of successful cases we observed significant variation in the finger movements due to a lack of precise reproducibility in the operation of the finger joint tendons. This led to a variable number of fingers making good contact with the lid circumference during the movement and a correspondingly varying "gear ratio" between the finger rotation and the lid rotation. This "gear ratio" can be measured by the vision system and will in future experiments constitute a good candidate for a quantitative quality measure enabling a fine-tuning of the unscrewing motion. However, even without a quantitative evaluation of this measure so far we observed that the endangerment of proper task execution by the unwanted variations in finger positions was strongly

counterbalanced by the availability of a redundant set of five fingers, ensuring that in most cases there existed sufficiently many "successful" finger contacts to drive the lid rotation.

When the lid was unscrewed, the system executed a manually triggered, preprogrammed lifting motion to remove the lid. One could use the achieved turning angle as a simple heuristic to define the end of the unscrewing motion automatically. The tactile sensors in the finger tips of the hand are not sensitive enough to notice this event.

To quantify the improvement in utilizing the proposed normalization heuristic of Sec. IV, we conducted an experiment, systematically probing whether an object pose is feasible for grasping or not. There exist several sources of possible failure: the robot might collide with an object in the scene (particularly with the other robot arm), violate the security box limiting the workspace, or violate the joint angle limits. Particularly, the restriction to the narrow security box bounded by the frame bars shown in Fig. 7a complicates the task, because the robot cannot extend its elbow beyond this virtual boundary. To avoid the joint angle limits, we exploit the redundancy of the 7-DOF robot arm computing a weighted Jacobean pseudo inverse [27] to reduce motion of joints near their limits.

For the evaluation, the pose space was sampled equidistantly covering a large portion of the workspace. The rotation angle around the object's symmetry axis covers the whole circle in steps of 30 degrees. The tilting of the object was neglected for this experiment. Fig. 7 visualizes the grasping success using different colors to indicate that picking of the object failed (white), succeeded partly, i.e. only hand-over failed (yellow), or succeeded completely (green). Three levels of elevation $z \in [0, 15, 30]cm$ of the object above the table are indicated by the nesting of circles, where the largest represents the table surface ($z=0$) and the smallest is set to $z=30cm$.

Fig. 7b (top) shows the results for the original determination of the tool frame (not exploiting the rotational symmetry), which is only successful for a location-specific range of object orientations as indicated by many white segments. In contrast, the optimized grasping strategy succeeds in grasping at nearly all positions, independent of the object's orientation as indicated by the circular green and yellow regions shown in Fig. 7b (bottom). The locations where grasping failed, correspond to regions where the elbow collides with the mounting frame (upper left corner), the hand collides with the other arm (right column) or the joint limits are exceeded due to the elevation of the object (upper right region). As indicated by the yellow segments, hand-over of the jar to the other hand systematically fails for a specific object orientation (90 degrees). Here, relaxing the constraint of using a fixed *absolute* hand-over pose of the arms, but only constraining their *relative* pose, will enlarge the redundant space and thus prevent a collision of the elbow with the virtual security box, which causes this failure.

VII. DISCUSSION

We showed that the combination of biologically motivated movement strategies can successfully generate fairly complex movements involving multifingered fine manipulation of objects. We focused on the coordination of redundant degrees of freedom at the arm and finger levels, utilizing two complementary strategies: a heuristic derived from the observation of human arm movements to maximize the set of object poses that can be handled by the system, and a learning approach for extracting a manifold of finger synergies from human training data to carry out the fine-manipulation phase. As a benchmark task, we focused on the removal of a lid from a glass jar by a periodic unscrewing motion carried out in a bimanual setting involving two anthropomorphic Shadow hands.

The accompanying video [26] demonstrates the whole manipulation action, including 6D object localization, object grasping from the human operator, hand-over between hands, the unscrewing itself, and the final lifting of the cap.

Still, many interesting issues remain for future work. So far, we used predefined arm and hand postures. Since the tactile sensors of the Shadow Hand have recently matured to a level where tactile-driven adaptation of the grasp posture becomes feasible, we plan to integrate tactile feedback both during grasping and during the cap turning motion. This will open additional possibilities for online-tuning of the cap motion and for the transitions between the now preprogrammed movement phases. With this added flexibility, a logical next step will be the generalization of the unscrewing motion to a larger set of lid types and to other types of container opening motions. We already have shown, that the manifold approach to learn complex motion trajectories is feasible also in other contexts [28].

The grasping performance is already very good with the present approach. In fact, all of the performed grasp trials were initially successful. Only in a relatively few number of cases we observed the glass slipping away from the grip during subsequent arm motion. This was due to heavy *tremor*-like oscillations of the robot arm arising from operating the arms beyond their torque specification limits due to the weight of the Shadow Hand subsystem.

The use of optimized tool poses for object pick-up, exploiting the symmetry of the object, significantly enhanced the feasible space for object grasping. Additionally, this strategy avoids awkward arm configurations. Extending the same method to compute a flexible hand-over pose for both robot arms, only constraining their relative pose, might eliminate the remaining hand-over problems.

ACKNOWLEDGMENTS

The authors acknowledge support from the DFG-funded CRC 673 – "Alignment in Communication", and the CoE 277 "Cognitive Interaction Technology" (CITEC). Further on, the authors thank Matthias Zenker for his technical support and his help with the video production.

REFERENCES

- [1] C. Borst, M. Fischer, and G. Hirzinger, "Calculating hand configurations for precision and pinch grasps," in *Proc. IROS*, 2002.
- [2] A. Miller, S. Knoop, H. Christensen, and P. Allen, "Automatic grasp planning using shape primitives," in *ICRA*, vol. 2, 2003.
- [3] Z. Xue, A. Kasper, M. Zoellner, and R. Dillmann, "An Automatic Grasp Planning System for Service Robots," in *ICAR*, 2009.
- [4] L. Natale and E. Torres-Jara, "A sensitive approach to grasping," in *6th Int. Workshop on Epigenetic Robotics*, 2006, pp. 87–94.
- [5] N. Pollard and V. Zordan, "Physically based grasping control from example," in *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 2005.
- [6] F. Röthling, R. Haschke, J. J. Steil, and H. Ritter, "Platform Portable Anthropomorphic Grasping with the Bielefeld 20-DOF Shadow and 9-DOF TUM Hand," in *Proc. IROS*, 2007.
- [7] J. Steffen, R. Haschke, and H. Ritter, "Experience-based and Tactile-driven Dynamic Grasp Control," in *Proc. IROS*, 2007.
- [8] A. Tsoli and O. C. Jenkins, "2D subspaces for user-driven robot grasping," in *RSS Workshop on Robot Manipulation: Sensing and Adapting to the Real World*, 2007.
- [9] M. Ciocarlie, C. Goldfeder, and P. Allen, "Dimensionality reduction for hand-independent dexterous robotic grasping," in *Proc. IROS*, 2007.
- [10] R. Platt Jr., A. Fagg, and R. Grupen, "Manipulation gaits: sequences of grasp control tasks," in *Proc. ICRA*, 2004.
- [11] D. Bentivegna, C. Atkeson, and G. Cheng, "Learning tasks from observation and practice," *Robotics and Auton. Systems*, vol. 47, 2004.
- [12] S. Schaal, A. Ijspeert, and A. Billard, "Computational approaches to motor learning by imitation," *Philosophical Transaction of The Royal Society of London: Ser.B, Biological Sci.*, vol. 358, no. 1431, 2003.
- [13] J. Steffen, S. Klanke, S. Vijayakumar, and H. Ritter, "Realising dextrous manipulation with structured manifolds using unsupervised kernel regression with structural hints," in *ICRA 2009 Workshop: Approaches to Sensorimotor Learning on Humanoid Robots*, 2009.
- [14] J. Peltason, I. Lütkebohle, B. Wrede, and M. Hanheide, "Mixed initiative in interactive robotic learning," in *Workshop on Improving HRI with Mixed-Initiative and Context-Awareness*, Toyama, 2009.
- [15] C. Borst, T. Wimböck, F. Schmidt, M. Fuchs, B. Brunner, F. Zacharias, P. R. Giordano, R. Konietzschke, W. Sepp, S. Fuchs, C. Rink, A. Albuschäffer, and G. Hirzinger, "Rollin' Justin - Mobile Platform with Variable Base," in *Proc. ICRA*, 2009, (Video submission).
- [16] H. J. Ritter, R. Haschke, and J. J. Steil, *A Dual Interaction Perspective for Robot Cognition: Grasping as a "Rosetta Stone"*, ser. Studies in Computational Intelligence. Springer, 2007, no. 77, pp. 159–178.
- [17] C. Bauchhage, S. Wachsmuth, M. Hanheide, S. Wrede, G. Sagerer, and H. Ritter, "The visual active memory perspective on integrated recognition systems," *Image and Vision Computing*, vol. 26, 2008.
- [18] I. Lütkebohle, J. Peltason, L. Schillingmann, C. Elbrechter, B. Wrede, S. Wachsmuth, and R. Haschke, "The curious robot - structuring interactive robot learning," in *ICRA*, 2009.
- [19] Shadow Robot Company, "The Shadow Dextrous Hand." [Online]. Available: <http://www.shadowrobot.com/hand/overview.shtml>
- [20] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics & Automation Magazine*, vol. 11, no. 4, 2004.
- [21] M. Cutkosky and R. D. Howe, "Human grasp choice and robotic grasp analysis," in *Dextrous Robot Hands*, S. T. Venkataraman and T. Iberall, Eds. Springer, 1990.
- [22] J. Steffen, R. Haschke, and H. Ritter, "Towards Dextrous Manipulation Using Manifolds," in *Proc. IROS*, 2008.
- [23] J. Steffen, S. Klanke, S. Vijayakumar, and H. Ritter, "Towards Semi-supervised Manifold Learning: UKR with Structural Hints," in *Proc. WSOM*, 2009.
- [24] P. Meinicke, S. Klanke, R. Memisevic, and H. Ritter, "Principal Surfaces from Unsupervised Kernel Regression," *TPAMI*, vol. 27, no. 9, 2005.
- [25] E. Oztop, L.-H. Lin, M. Kawato, and G. Cheng, "Dexterous Skills Transfer by Extending Human Body Schema to a Robotic Hand," in *Proc. Humanoids*, 2006.
- [26] R. Haschke, J. Steffen, C. Elbrechter, and H. Ritter, "accompanying video for this paper." [Online]. Available: <http://www.youtube.com/user/citecbielefeld#p/u/8/XQ5G0mKoDfc>
- [27] T. Chan and R. Dubey, "A weighted least-norm solution based scheme for avoiding joint limits for redundant joint manipulators," *IEEE Trans. on Robotics & Automation*, vol. 11, no. 2, pp. 286–292, Apr 1995.
- [28] J. Steffen, M. Pardowitz, and H. J. Ritter, "Using structured ukr manifolds for motion classification and segmentation," in *IROS*, 2009, pp. 4785 – 4790.