Human-Inspired Selection of Grasp Hypotheses for Execution on a Humanoid robot

Markus Przybylski
Tamim Asfour
Rüdiger Dillmann
Institute for Anthropomatics
Karlsruhe Institute of Technology
Karlsruhe, Germany
Email: {markus.przybylski,asfour,dillmann}@kit.edu

René Gilster
Heiner Deubel
Department Psychologie
Ludwig-Maximilians-Universität
Munich, Germany
Email: {rene.gilster,heiner.deubel}@psy.lmu.de

***Abstract*—Future humanoid robots will need the capability to grasp and manipulate arbitrary objects in order to assist people in their homes, to interact with them and with the environment. In this work, we present an approach to grasp known objects. Our approach consists of an offline step for grasp planning, a rating step which determines the human likeness of the grasps and an execution step, where the most suitable grasp is performed on a humanoid robot. We especially focus on the rating step where we use human grasping data to rate pre-computed grasp hypotheses from our grasp planner in order to select the most human-like feasible grasp for execution on the real robot. We present the details of our method together with experiments on our ARMAR-III humanoid robot.**

I. INTRODUCTION

The capability to grasp and manipulate objects is crucial for future service robots that will help people in their daily lives. In robotics, grasping of known objects consists of several steps. In the first step, a grasp planning method is used to generate a set of grasp hypotheses, given the object, the robotic hand, and possibly obstacles in the environment. In the second step a feasible grasp has to be selected for execution on the robot in the third step. Grasp selection is usually subject to several constraints. Grasps can be rated depending on their force-closure score [1] which indicates how well the grasps can resist external forces and torques. In real-world environments also obstacles have to be taken into account and the robot's embodiment can turn many grasp hypotheses infeasible, because they are not reachable due to kinematic constraints. Yet, with grasp planners able to generate hundreds or thousands of grasps, often many feasible grasps remain to select from, despite the constraints stated above. Also, the force-closure score does not tell whether a human being would choose a specific grasp or not. It seems desirable for humanoid robots that they act human-like. We are convinced that a humanoid robot that performs tasks and grasps objects the way a human would do will be more readily accepted by people than a robot grasping objects in unintuitive ways. We therefore believe that grasp selection by humanoid robots should be motivated by human grasping. This leads to the question how to decide which grasps are more human-like than others. In this paper we present an intuitive approach to rate grasps generated by our grasp planner, given measured data from the human grasping process.

II. RELATED WORK

Before we describe our approach in detail we give an overview of related work in the field. In the recent years, a multitude of humanoid robots have been presented, for example HRP2 ([2],[3]), ARMAR ([4],[5]), Dexter [6], Justin [7], or NASA Robonaut [8] where an important focus of research was grasping and manipulation, also including dealing with furniture and doors. Work has also been conducted on pre-grasp manipulation [9], where a flat object is slided to an intermediate position where it can be grasped more easily.

In the area of grasp planning for known objects, a number of simulation-based approaches were introduced based on simulation environments such as GraspIt! [10], OpenRAVE [11] and Simox [12]. In the context of the *grasping by parts* paradigm several authors proposed shape approximation techniques to prune the search space for grasp planning. The first of these were Miller et al. [13] who manually decomposed objects into basic shape primitives such as boxes, spheres, cylinders and cones. Goldfeder et al. [14] used superquadrics and Huebner et al. [15] proposed an approach to decompose objects into a set of minimum volume bounding boxes. Following the same line of thinking, but guided by the idea of improved shape approximation accuracy, we proposed in our earlier work the use of object representations based on the medial axis [16] and the medial axis transform [17]. Aleotti et al. [18] presented an approach for grasp planning based on the Reeb Graph. Vahrenkamp et al. [19] proposed a method that integrates grasp planning with searching for collision free grasp motions.

As for the human component in grasp selection, recently a human-guided grasp measure was introduced by Balasubramanian et al. [20]. In their experiments the authors let test persons interact haptically with a robot so the robot arm would grasp the objects in a way the test persons considered intuitive. The authors identified orthogonality of wrist orientation towards the object's principal axis as a measure that is optimized by humans when selecting grasps for robots.

As to human grasping, a variety of behavioral studies have addressed the question of how hand posture and finger contact positions are selected when objects are grasped. These studies have demonstrated that how an object is grasped depends on a number of parameters. Paulignan et al. [21] showed that cylindrical objects that were placed at different positions in front of the participant were grasped at different positions on the object. The center of mass of the object is also known to influence the choice of contact points [22]. Finally, the intended goals of an action or planned movement sequences also influence how objects are grasped ([23],[24]). Notably, the choice of fingertip contact points has been studied in most cases only for two-finger grasps that involved the thumb and index finger. This contrasts with the recent finding that the *default grasp* of humans seems to involve all five fingers [25]. In general, the small variability of contact points and the preshaping of the hand in early phases of the prehension movement indicate a preplanning in the choice of contact points [26]. This is in line with the view that the goal posture of the hand in grasping is specified in advance, followed by an appropriate movement to this posture [27].

The remainder of this paper is organized as follows. In section III we describe the components of our system. In section IV we explain in detail our experiments and evaluate the results before we conclude the paper in section V with ideas for further research.

III. OUR APPROACH

In this section we explain our approach in detail. Our method consists of four steps: The generation of grasp hypotheses by a grasp planner (see section III-A), the process of collecting human grasping data (III-B), the rating of the human-likeness of the grasps (III-C) and finally the selection and execution of a grasp on the real robot (III-D).

A. Grasp Hypotheses Generation

For grasp hypotheses generation we use a technique based on a novel object representation we recently presented in [28]. Our object representation, the *grid of medial spheres*, is based on the medial axis transform (MAT) [29], a complete shape descpritor, which can represent arbitrarily shaped objects with high accuracy and additionally contains volumetric information of the object. The MAT can be obtained by inscribing spheres of *maximum diameter* into the original shape, where *maximum diameter* means that each inscribed sphere has to touch the original shape at two distinct points. The sphere diameter d and the object angle α_o [30], which describes the maximum angle included by two vectors from the sphere's center to two different nearest surface points of the original shape touched by this sphere, are parameters that define the level of detail of the object representation. By discarding spheres with small object angles α_{o_i} and small diameters d_i we can eliminate edges and corners of the object.

Our goal is to generate grasps where the thumb opposes the other fingers. This kind of grasp hypotheses directly emerges from our *grid of medial spheres* object representation. We use principal component analysis of local areas of the inscribed spheres' centers to estimate local symmetry planes and axes of the object. These local symmetry properties, combined with the diameters d_i of the associated spheres contain valuable information how the object can be grasped. Approach directions and orientations of the hand are generated in a way that the robot hand is able to wrap around the local symmetry axis or that the thumb and the other fingers touch the object at opposing sides in case of a local symmetry plane. For more details on this method we refer the reader to [28]. The grasp hypotheses generated by this process are tested for force-closure. The force-closure grasps among these grasp hypotheses will later be rated according to their similarity to human grasps. We will describe this in detail in section III-C.

B. Human Grasping Data Collection Method

We collect human grasping data at a sampling rate of 240 Hz using a Polhemus Liberty electromagnetic motion tracking system, with a static accuracy of 0.8 mm RMS for the Cartesian x, y, and z positions and 0.15 degrees for the orientation of the sensor markers. We attach the sensors with medical tape to the fingernails of all five fingers of the right hand. We use an additional wrist sensor as a reference for the computation of parameters for the transport phase of the hand. The thickness of the individual fingers and the orientation of the sensor markers are taken into account in order to compute the actual fingertip contact positions on the objects. During the data collection session, participants wear liquid-crystal shutter glasses [31] that can rapidly switch from opaque to transparent to ensure standardized viewing conditions. The whole procedure is controlled by custom made software running on a PC. We filter the obtained position coordinates with a second-order Butterworth filter that employs a low-pass cut-off frequency of 15 Hz. Velocities of each finger are determined from the data stream. We determine the start and end of the movement with a velocity threshold criterion of 0.1 m/s for the wrist sensor. Contact points of the fingertips are determined in two steps: First, the sample at which the velocity of the wrist marker is lower than 0.1 m/s is determined. We then define the minimum velocity within the next 40 frames as the end of the transport phase of the hand. Potential contact times for individual fingers are calculated separately. They are determined by searching for the minimum velocity within the next 40 frames, starting from the end of the transport phase. Second, we use a spatial criterion to decide if the fingers make contact with the object. If the distance between the finger position and the object is smaller than 0.5 cm, we assume that the finger touches the object.

C. Rating the Human-Likeness of Grasps

We now face the problem of deciding which grasp hypotheses generated by our grasp planner are more human-like than others. Each human grasp h from the data recorded as described in section III-B contains the final fingertip position p_{h_i} on the object for each finger i. Each force-closure grasp g

from our grasp planner consists of the final wrist pose p_w of the hand and the final joint angle vector q of the finger joints:

$$g = (p_w, q) \tag{1}$$

In order to compute the fingertip position p_{g_i} of finger i of a grasp g from wrist pose p_w and joint angle vector q we use the forward kinematics model $FwdKin()$ of our simulated robot hand:

$$p_{g_i} = FwdKin(p_w, q_i) \tag{2}$$

We define a criterion to rate similarity between a planned grasp g and a human grasp h:

$$d_z = \sqrt{\sum_{i=1}^{5}(p_{g_{i,z}} - p_{avg(h)_{i,z}})^2} \tag{3}$$

In the equation above $p_{g_{i,z}}$ denotes the z component of the i-th fingertip of the grasp g planned by the grasp planner. $p_{avg(h)_{i,z}}$ denotes the z component of the i-th fingertip postion of the average of all grasps h of one human test subject. The z axis of our object coordinate frame points up. The criterion in equation (3) does not compute the overall displacement of fingertips, but only their displacement in z direction, i.e. if the fingertips of the planned grasp touch the object at about the same height as the human grasp. We will explain the reason for this choice in more detail later in this paper.

D. Selecting and Executing Grasp Hypotheses

We now have a set of force-closure grasps from the grasp planner, which have been rated with the criterion d_z introduced in the previous section. Which of these grasps should now be executed on the real robot? The grasps with the lowest values of d_z are the grasps most similar to the human grasps.

In practice however, the feasibility of grasps depends on the actual pose of the object in the scene. Due to kinematic constraints of the robot and obstacles in the scene, not all precomputed grasps are reachable in a given situation. In order to select a feasible grasp for execution, the object has first to be identified and its pose has to be estimated. We use an approach for textured objects described in [32]. First the object is recognized using 2D feature correspondences. Then 2D localization using scale-invariant feature transform (SIFT) descriptor correspondences is performed. Finally, a 6D pose estimate of the object is calculated based on the outcome of the 2D localization.

As now the object pose is known, the grasps which are originally in the object coordinate frame can be transformed to the robot's platform coordinate system. We order the grasps with respect to their d_z rating, the hypothesis with the smallest rating first. For each grasp we check reachability using inverse kinematics, where all seven joints of ARMAR's right arm, as well as the hip joint, are considered. As soon as we find a grasp which is actually reachable, we select it for execution. To actually perform the selected grasp on ARMAR, we use a visual servoing approach [33]. For this purpose we first generate a pre-pose of the hand by shifting the final grasp pose for a small distance along the negative approach direction. This is necessary to make sure that the hand approaches the final grasp pose from the desired direction without toppling the object. During the grasping movement, ARMAR moves its hand first to the pre-pose of the selected grasp, then to the final grasp pose. At the final grasp pose the hand is closed and ARMAR lifts the object.

IV. EXPERIMENTS

In this section we present our experiments which comprise three parts: The collection and evaluation of human grasping data (IV-A), the generation of grasps and their rating with respect to the human data (IV-B), and the selection and execution of grasps on the robot (IV-C).

A. Collection and Evaluation of Human Grasping Data

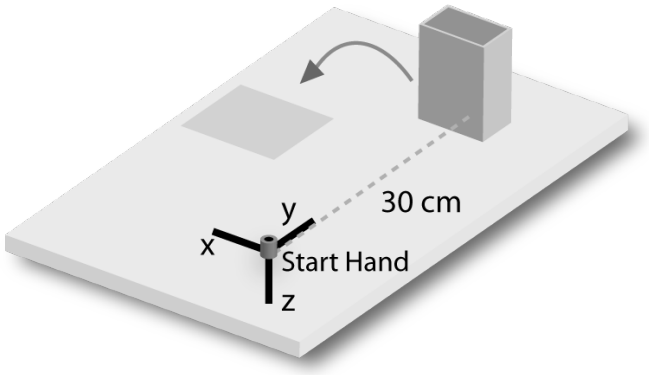

Fig. 1: Experimental setup for collecting human grasps. The object was to be picked up and placed in the grey shaded area.

The experimental setup for collecting human grasping data is shown in Fig. 1. We used a cylinder (6.3 cm diameter, 12.2 cm height) and a rectangular box (8.4 cm $\times$ 5.2 cm $\times$ 15.4 cm) as test objects. The distance of the objects (from their center of mass) to the starting position of the hand was 30 cm in the y direction. The 'long' side of the box was always parallel to the x axis. At the beginning of each trial, thumb and index finger of the right hand touched a small knob (starting position) which also served as the origin of the coordinate system. Participants looked at the object for 1000 ms after the opening of the shutter glasses. A tone (800 Hz, 100 ms) served as the signal to start the grasping movement. We instructed the participants to grasp the object 'from the side' with as many fingers as they liked in their self chosen natural speed. No instructions were given with respect to contact points of the fingers with the object. The object was to be picked up and placed in the target region (grey shaded area in Fig. 1). The hand was brought back to the starting position and the shutter glasses became opaque again. Each object was grasped ten times in random order. We investigated five right handed participants. Results from two exemplary participants are shown in Fig. 2. The left column depicts the contact positions with the object and the right column illustrates the

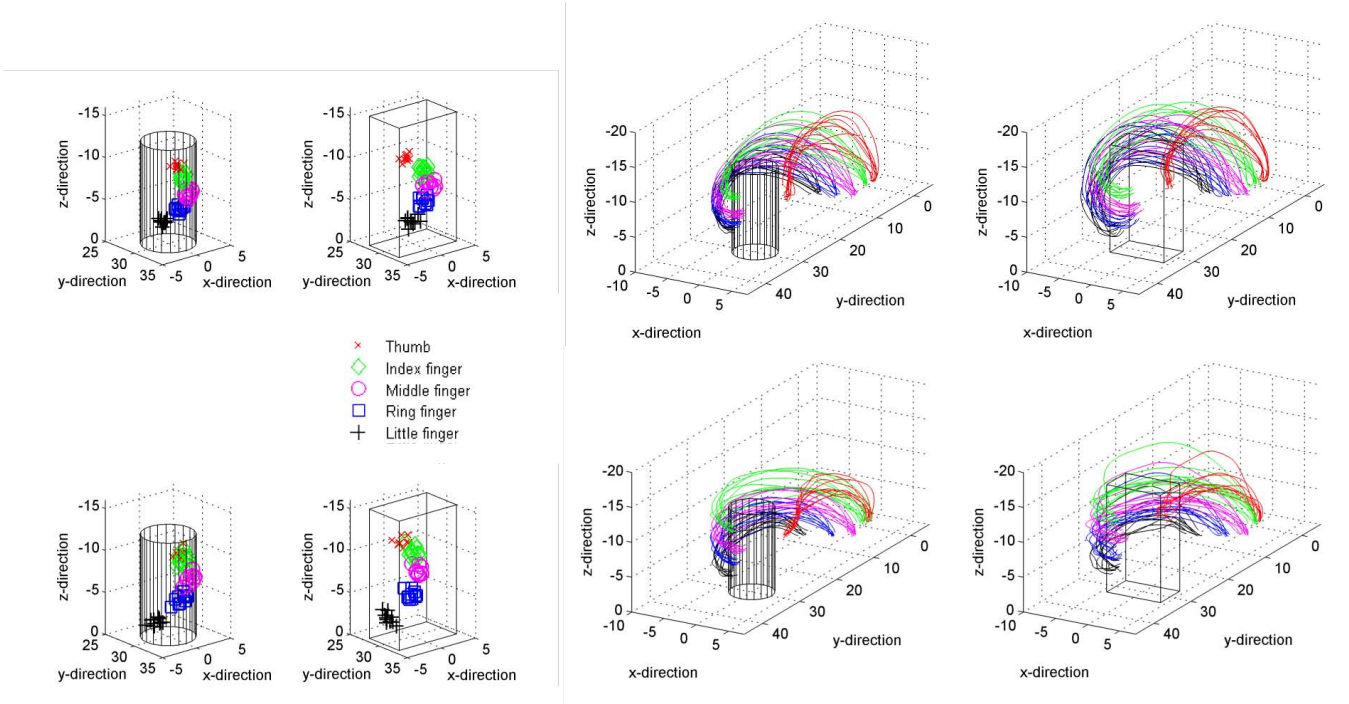

Fig. 2: Final fingertip positions (left column) and trajectories of individual fingers (right column) for two exemplary participants (rows).

trajectories of individual fingers until contact is made with the object. It can be seen that participants were quite consistent in their choice of contact points, although differences exist between participants for this choice. For example, in the box conditions the participant in the upper row placed all fingers on the long side of the box, while the participant shown in the lower row placed the little finger always on the small side.

B. Planning and Rating Grasps

For the grasp planning, rating and execution parts of our experiments we used the following two objects: A cylindric salt can and a medium-sized cereals box. The model of the cylindric salt can was obtained using a 3D laser scanner ([34],[35]). The cereals box was modeled by hand. For rating the grasps from the grasp planner, we chose to use the data from a single participant since averaging final fingertip positions between participants will not lead to meaningful results in most instances. For this work, we randomly chose to use the data shown in the upper row in Fig. 2. In order to make the human grasping data from the previous step applicable to the differing sizes of the test objects used with the robot, we scaled up the human grasping data to fit the size of these test objects. Scaling of the grasping data was performed with respect to the object's center of gravity, proportionally to the scaling of the object size. For the work in this paper, we assumed that all objects have graspable sizes, i.e. scaling the human data would not lead to infeasible grasps.

Figure 3 shows the results of the grasp planning and rating process. Green and orange rays towards the object describe approach directions of the hand. Magenta lines at the end of the approach directions indicate the respective hand orientations which can be thought of as an imaginary axis the fingers of the hand wrap around when they close starting from a parallel preshape. For the cereals box, our grasp planner generated grasps where the approach directions and hand orientations are inside the object's biggest symmetry plane. This ensures that the hand can close around the object with the thumb touching the object at the side opposed to the side touched by the other fingers. In case of the salt can, the grasp planner exploited the object's central symmetry axis and generated grasps with approach directions perpendicular to the symmetry axis and hand orientations parallel to the symmetry axis. All the depicted grasps are force-closure grasps. The black spheres indicate the fingertip positions from the human grasping experiments, transformed to the object coordinate frame and scaled to fit the size of the object. The purple spheres located at the approach directions describe the final wrist positions of the hand during grasp planning. The sizes of these spheres are inverse to the value of the d_z rating introduced in section 3, i.e. the biggest spheres represent the grasps with the smallest d_z rating, that is, the grasps with the biggest similarity to the data from the human grasping experiments. We note that for both objects the planned grasps most similar to the human grasps with respect to the d_z rating are located somewhere between the top and the center of mass of the objects, just as we would anticipate from the human grasping experiments.

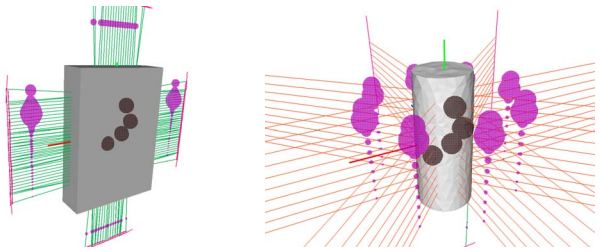

Fig. 3: Results from the grasp planning and rating procedure for the cereals box (left) and the cylindric salt can (right). Green and orange rays describe approach directions of the hand towards the object. Magenta lines at the end of the approach directions indicate the respective hand orientation vectors. Black spheres indicate the fingertip positions from the human grasping experiments. Purple spheres located at the approach directions describe the final wrist positions of the hand during grasp planning. The bigger these spheres, the bigger the similarity of the respective grasps to the data from the human grasping experiments based on the proposed similarity criterion.

C. Selecting and Executing Grasps

As already described in section III-D, the grasps generated by the grasp planner are sequentially tested for kinematic reachability, starting with the most human-like grasp. The first grasp actually reachable is automatically selected and performed on ARMAR-III. Figures 4 and 5 show the actual grasping process on our humanoid robot ARMAR-III for the two test objects. The first row of Fig. 4 shows ARMAR localizing the respective object on a table in front of it. In the second row the object pose is updated in the simulation environment. In the third row all grasps and pre-poses generated by the grasp planner for the object are decpicted. The fourth row shows the selection of the first actually reachable grasp for execution. The first row of Fig. 5 shows ARMAR's hand reaching the pre-pose (first row), then the final grasp pose (second row), closing the hand (third row) and finally lifting the object (fourth row).

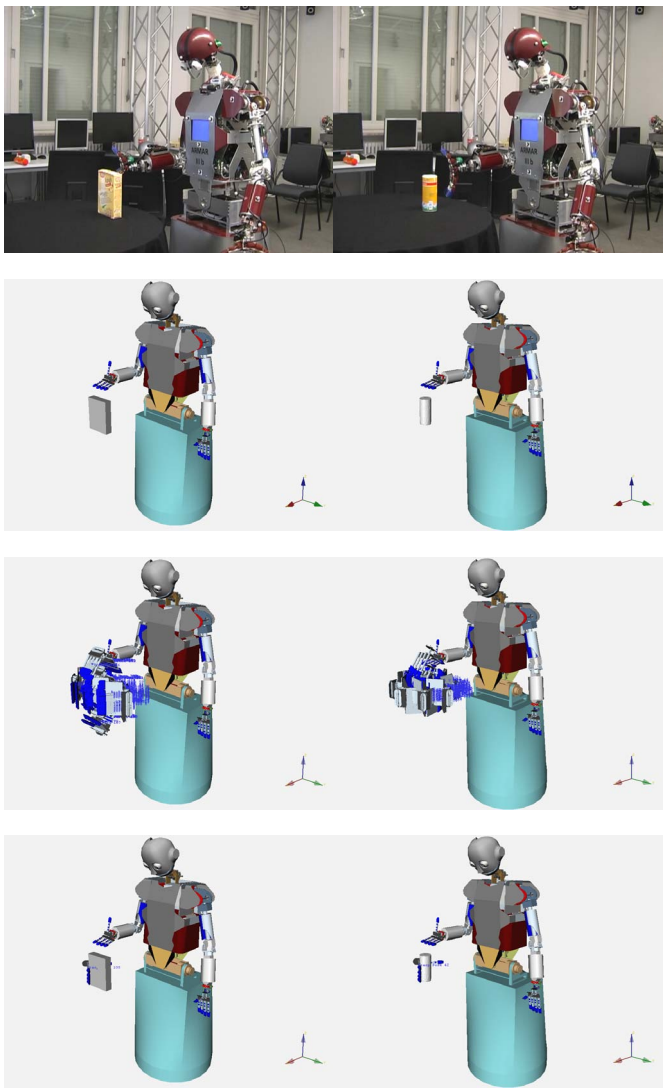

Fig. 4: Grasp rating and selection. Object localization (first row), updating the object pose in the simulation environment (second row), all grasps generated by the grasp planner and rated with respect to human grasping data (third row), selecting the first feasible grasp for execution (fourth row).

V. DISSCUSSION AND CONCLUSION

In this paper we proposed an approach to grasp known objects with a humanoid robot, with a focus on human-inspired selection of the grasp actually to be performed. The results show that human grasping data can be used to rate grasps generated by a grasp planner with respect to their human-likeness, so the selection of grasps for execution on a humanoid robot can be biased towards more human-like grasps. The results also show that our grasp planning algorithm produces grasps similar to those intuitively used by human beings. The human grasping data seem to indicate that object symmetries, as exploited by our grasp planner, affect humans in their choice how to grasp an object.

Yet we are aware of the limitations of our rating approach. One possible point of criticism might be our choice of similarity criterion that only considers deviations along the z direction of the object coordinate frame, i.e. variations of the height where the object is grasped affect the rating, while variations of the directions, from which side the object is grasped, do not. One might argue that rather a criterion based on the complete euclidean distance should be used. During our experiments, we actually considered this option, too, but found that in the case of our highly symmetrical test objects, especially the cylindric salt can, a euclidean distance criterion would be too restrictive, for the following reason: The grasps generated by the grasp planner relate to the object coordinate frame, which does not respect the object's symmetry. If we turn the cylindric object around its symmetry axis on the table in front of the robot, then the object pose changes and the most human-like grasps with respect to a euclidean distance criterion may be located

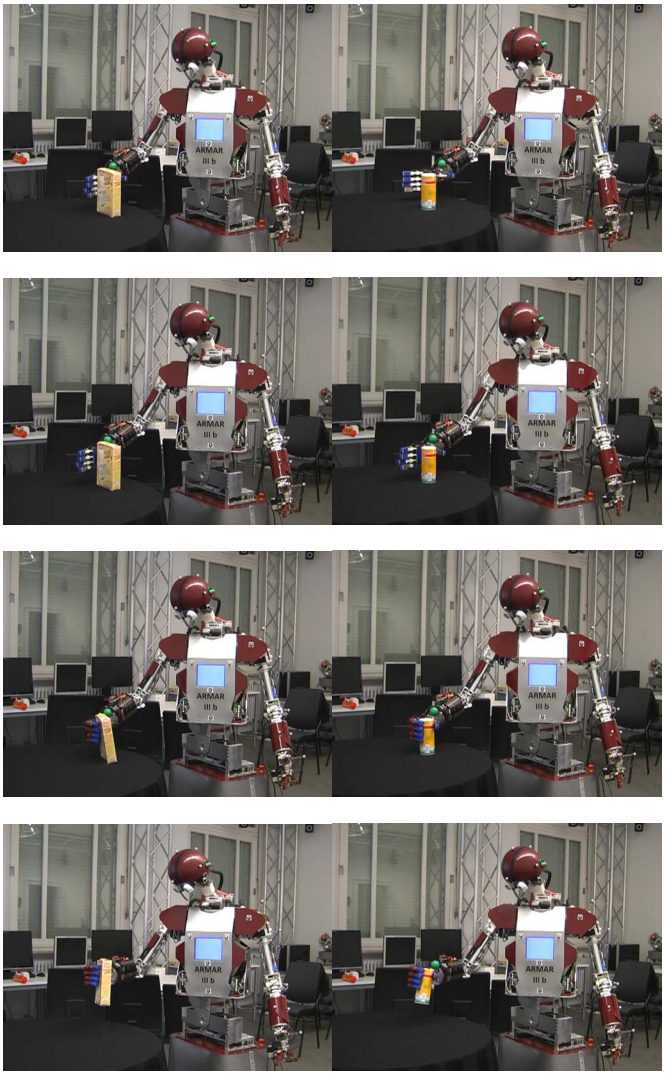

Fig. 5: Grasp execution on ARMAR-III. The hand at the prepose (first row), at the grasp pose (second row), closing the hand (third row), and lifting the object (fourth row).

at the back of the object, unreachable for the robot.

Apart from this symmetry issue, our experiments raised a set of interesting questions. First, we only used two objects with relatively simple shapes - a box and a cylinder - for the data collection experiments and for grasping on the robot. While this is motivated by pychological findings and the collected human data can be scaled up or down to fit objects of similar shapes but different sizes, the question remains how to deal with more complexly shaped objects. One might think of grasps at the handle of a pitcher, at the rim of a salad bowl. Or consider the classic coffee cup: It can be grasped from the top, from various sides, at the handle or at the rim with some fingers reaching inside the cup. This leads us to the second question: What kind of data can we expect to get from the human grasping experiments, if we do not tell the test subjects how to grasp the object? In our experiments in this paper, we explicitly told the test subjects to grasp the objects from the side. The data for each of the subjects show little variation regarding the final fingertip positions during grasping (see Fig. 2). If subjects are asked to grasp the test objects as they please, and if the test objects have more complex shapes, this could lead to much more variation in the recorded human grasping data. In our opinion there are two conceivable outcomes: Either the recorded data contain distinct clusters or not. In the case of clusters, one might think of similar grasps belonging to the same cluster: one cluster for grasps at the handle of a cup, another cluster for grasps from the top, and so on. In this case one might think of using our similarity criterion with different clusters as human reference data, depending on the respective way the robot is asked to grasp the object. For instance one would use only the human grasps of the *handle grasps* cluster as a reference for rating robot grasps, if the robot is asked to grasp the cup at the handle. The more challenging case might be the presence of very strangely shaped clusters in the data or the absence of clusters at all. In these cases there would be considerable variation in the data, which leads us to the third question: How can a criterion for human-likeness deal with big variations in the collected human grasping data? In the experiments of this paper our approach of calculating a distance criterion between planned grasps and the average of the observed grasping data of a single human subject was based on the assumption that averaging human grasps would yield a feasible *average grasp*. Yet this might not always be the case. Consider a big salad bowl for example, where one-handed grasps can only be performed at the rim, but the test subject might grasp the rim at very differenct places during the experiments. In this case averaging seems very questionable, since computing the average over all observed grasps would possibly not yield a valid grasp at all. We believe that in these cases the usage of machine learning methods might be reasonable, where the human-likeness of a planned query grasp could be rated depending on whether there exists *any* human grasp in the space of observed human grasps for this object within a given distance to the query grasp. We consider the questions raised above to be promising topics of future research.

ACKNOWLEDGMENT

This work was supported by EU through the project GRASP.

REFERENCES

[1] C. Ferrari and J. Canny, "Planning optimal grasps," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*, May 1992, pp. 2290 –2295 vol.3.

[2] K. Okada, T. Ogura, A. Haneda, J. Fujimoto, F. Gravot, and M. Inaba, "Humanoid motion generation system on hrp2-jsk for daily life environment," in *Mechatronics and Automation, 2005 IEEE International Conference*, vol. 4, july-1 aug. 2005, pp. 1772 – 1777 Vol. 4.

[3] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner, "Grasp planning in complex scenes," in *Humanoid Robots, 2007 7th IEEE-RAS International Conference on*, 29 2007-Dec. 1 2007, pp. 42 –48.

[4] T. Asfour, K. Regenstein, P. Azad, J. Schrder, A. Bierbaum, N. Vahrenkamp, and R. Dillmann, "Armar-iii: An integrated humanoid platform for sensory-motor control," in *Proceedings of the 7th IEEE-RAS International Conference on Humanoid Robots*, 2006.

[5] T. Asfour, P. Azad, N. Vahrenkamp, K. Regenstein, A. Bierbaum, K. Welke, J. Schröder, and R. Dillmann, "Toward Humanoid Manipulation in Human-Centred Environments," *Robotics and Autonomous Systems*, vol. 56, pp. 54–65, January 2008.

[6] R. Platt, "Learning and generalizing control based grasping and manipulation skills," Ph.D. dissertation, Department of Computer Science, University of Massachusetts Amherst, 2006.

[7] T. Wimbock, C. Ott, and G. Hirzinger, "Impedance behaviors for two-handed manipulation: Design and experiments," in *Robotics and Automation, 2007 IEEE International Conference on*, april 2007, pp. 4182 –4189.

[8] T. Martin, R. Ambrose, M. Diftler, J. Platt, R., and M. Butzer, "Tactile gloves for autonomous grasping with the nasa/darpa robonaut," in *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, vol. 2, 26-may 1, 2004, pp. 1713 – 1718 Vol.2.

[9] D. Kappler, L. Chang, M. Przybylski, N. Pollard, T. Asfour, and R. Dillmann, "Representation of Pre-Grasp Strategies for Object Manipulation," Nashville, USA, December 2010.

[10] A. T. Miller and P. K. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics Automation Magazine, IEEE*, vol. 11, no. 4, pp. 110 – 122, Dec. 2004.

[11] R. Diankov and J. Kuffner, "Openrave: A planning architecture for autonomous robotics," Robotics Institute, Tech. Rep. CMU-RI-TR-08-34, July 2008. [Online]. Available: http://openrave.programmingvision.com

[12] [Online]. Available: http://www.sourceforge.net/projects/simox

[13] A. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic grasp planning using shape primitives," in *Robotics and Automation, 2003. Proceedings. ICRA '03. IEEE International Conference on*, vol. 2, Sept. 2003, pp. 1824 – 1829 vol.2.

[14] C. Goldfeder, C. Lackner, R. Pelossof, and P. K. Allen, "Grasp planning via decomposition trees," in *Robotics and Automation, 2007 IEEE International Conference on*, April 2007, pp. 4679 –4684.

[15] K. Huebner, S. Ruthotto, and D. Kragic, "Minimum volume bounding box decomposition for shape approximation in robot grasping," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, May 2008, pp. 1628 –1633.

[16] M. Przybylski, T. Asfour, and R. Dillmann, "Unions of balls for shape approximation in robot grasping," in *Intelligent Robots and Systems, 2010. IROS 2010. IEEE/RSJ International Conference on*, 2010.

[17] ——, "Planning grasps for robotic hands using a novel object representation based on the medial axis transform," in *Intelligent Robots and Systems, 2011. IROS 2011. IEEE/RSJ International Conference on*, 2011.

[18] J. Aleotti and S. Caselli, "Grasp synthesis by 3d shape segmentation using reeb graphs," in *Intelligent Robots and Systems, 2010. IROS 2010. IEEE/RSJ International Conference on*, 2010.

[19] N. Vahrenkamp, M. Do, T. Asfour, and R. Dillmann, "Integrated grasp and motion planning," in *ICRA*, Anchorage, USA, Mai 2010.

[20] R. Balasubramanian, L. Xu, P. Brook, J. Smith, and Y. Matsuoka, "Human-guided grasp measures improve grasp robustness on physical robot," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, may 2010, pp. 2294 –2301.

[21] Y. Paulignan, V. G. Frak, I. Toni, and M. Jeannerod, "Influence of object position and size on human prehension movements," *Experimental Brain Research*, vol. 114, no. 2, pp. 226–234, 1997.

[22] G. P. Bingham and M. M. Muchisky, "Center of mass perception and inertial frames of reference," *Perception & Psychophysics*, vol. 54, no. 5, pp. 617–632, 1993.

[23] C. Ansuini, L. Giosa, L. Turella, G. Altoe, and U. Castiello, "An object for an action, the same object for other actions: effects on hand shaping," *Experimental Brain Research*, vol. 185, no. 1, pp. 111–119, 2008.

[24] C. Hesse and H. Deubel, "Advance planning in sequential pick-and-place tasks," *Journal of NeuroPhysiology*, vol. 104, no. 1, pp. 508–516, 2010.

[25] R. Gilster, C. Hesse, and H. Deubel, "Contact points during multidigit grasping of geometric objects," submitted.

[26] L. Schettino, S. V. Adamovich, and H. Poizner, "Effects of object shape and visual feedback on hand configuration during grasping," *Experimental Brain Research*, vol. 151, no. 2, pp. 158–166, 2003.

[27] D. A. Rosenbaum, R. J. Meulenbroek, J. Vaughan, and C. Jansen, "Posture-based motion planning: Applications to grasping," *Psychological Review*, vol. 108, no. 4, pp. 709–734, 2001.

[28] M. Przybylski, T. Asfour, and R. Dillmann, "Planning grasps for robotic hands using a novel object representation based on the medial axis transform," iROS 2011.

[29] H. Blum, *Models for the Perception of Speech and Visual Form*. Cambridge, Massachusetts: MIT Press, 1967, ch. A transformation for extracting new descriptors of shape, pp. 362–380.

[30] B. Miklos, J. Giesen, and M. Pauly, "Discrete scale axis representations for 3d geometry," in *ACM SIGGRAPH 2010 papers*, ser. SIGGRAPH '10. New York, NY, USA: ACM, 2010, pp. 101:1–101:10.

[31] P. Milgram, "A spectacle-mounted liquid-crystal tachistoscope," *Behavior Research Methods, Instruments & Computers*, vol. 19, no. 5, pp. 449–456, 1987.

[32] P. Azad, T. Asfour, and R. Dillmann, "Stereo-based 6d object localization for grasping with humanoid robot systems," in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, 29 2007-nov. 2 2007, pp. 919 –924.

[33] N. Vahrenkamp, S. Wieland, P. Azad, D. Gonzalez, T. Asfour, and R. Dillmann, "Visual servoing for humanoid grasping and manipulation tasks," in *Humanoid Robots, 2008. Humanoids 2008. 8th IEEE-RAS International Conference on*, dec. 2008, pp. 406 –412.

[34] A. Kasper, R. Becher, P. Steinhaus, and R. Dillmann, "Developing and analyzing intuitive modes for interactive object modeling," in *International Conference on Multimodal Interfaces*, 2007.

[35] R. Becher, P. Steinhaus, R. Zöllner, and R. Dillmann, "Design and implementation of an interactive object modelling system," in *International Symposium on Robotics*, 2006.