Modeling and Evaluation of Human-to-Robot Mapping of Grasps

Javier Romero Hedvig Kjellström Danica Kragic

***Abstract*— We study the problem of human to robot grasp mapping as a basic building block of a learning by imitation system. The human hand posture, including both the grasp type and hand orientation, is first classified based on a single image and mapped to a specific robot hand. A metric for the evaluation based on the notion of virtual fingers is proposed. The first part of the experimental evaluation, performed in simulation, shows how the differences in the embodiment between human and robotic hand affect the grasp strategy. The second part, performed with a robotic system, demonstrates the feasibility of the proposed methodology in realistic applications.**

I. Introduction

Programming service robots for new tasks puts significant requirements on the programming interface and the user. Programming by Demonstration (PbD) systems offer a great opportunity to unexperienced users for integrating complex tasks in a robotic system. However, representing, detecting and understanding human activities has been proven difficult and has been investigated closely during the past several years, [1], [2], [3], [4].

In the past, we have studied imitation of object manipulation tasks, using magnetic trackers, [4]. Although magnetic trackers and datagloves deliver exact values of hand joints, it is desirable that the user demonstrates tasks to the robot in a natural way; the use of gloves or other types of sensors may prevent a natural grasp. This motivates the use of systems based on visual input. In this paper, we concentrate on the use of our vision based grasp classification system presented in [5] for evaluation and execution of the grasp mapping on a robot. The contributions of the work presented here are:

- Based on a single image and the classification methodology [5], we purpose a mapping strategy between a human and two robot hands of different kinematical properties. The mapping is performed according to the grasp taxonomy proposed in [6], shown in Fig. 1.
- The distinction between the grasp categories is made based on the preshape of the hand and also in terms of different strategies for approaching the objects.
- We propose a metric for evaluation of the mapping strategy and use it in the experimental evaluation.

We start with the state of the art description in Section II, followed by a short overview of grasp classification in Section III. Section IV describes the grasp mapping strategy and Section V presents the evaluation of the system. The paper is concluded in Section VI.

Authors are with the Centre for Autonomous Systems, CVAP, KTH, Stockholm, Sweden. {jrgn,hedvig,dani}@ckth.se. The work was supported by the EU projects GRASP, IST-FP7-IP-215821 and PACO-PLUS, FP6-2004-IST-4-27657 and the Swedish Foundation for Strategic Research.

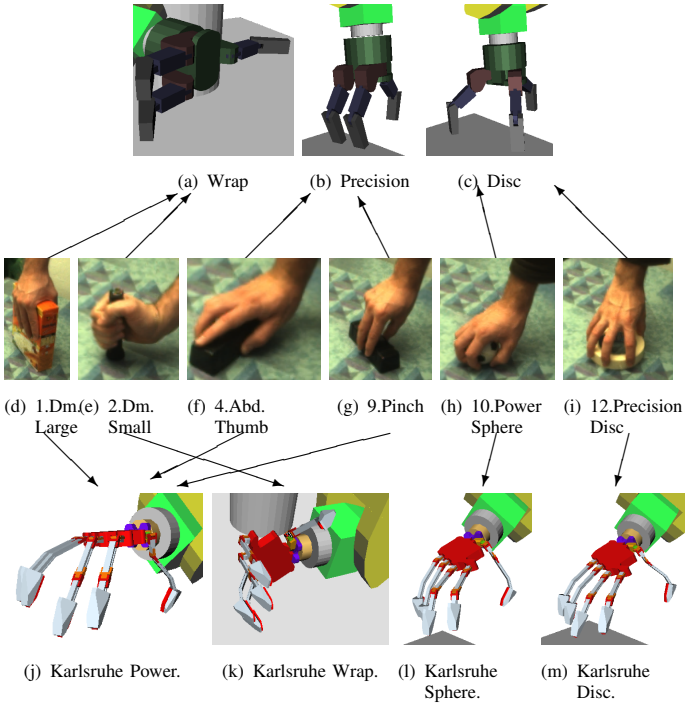

Fig. 1. The six grasps according to Cutkosky's grasp taxonomy [6] considered in the classification, and the three grasps for a Barrett hand and Karlsruhe hand, with human-robot class mappings ((d,e)→(a),(f,g)→(b), (h,i)→(c)), (d,f,g)→(j), (e)→(k), (h)→(l), (i)→(m) shown. a) Barrett Wrap. b) Barrett Precision. c) Barrett Precision Disc d) Large Diameter grasp, 1. e) Small Diameter grasp, 2. f) Abducted Thumb grasp, 4. g) Pinch grasp, 9. h) Power Sphere grasp, 10. i) Precision Disc grasp, 12. j) Karlsruhe Power, k) Karlsruhe Wrap, l) Karlsruhe Sphere, m) Karlsruhe Disc

II. Related Work

In the field of robotics, most of the object grasping systems are based on the a-priori object knowledge where either analytical or off-line methodology is used for grasp execution. In the work presented in this paper, we aim for learning grasps directly from human and mapping them on different robot hands. This relates to the work of [7] that classify objects based on their affordances (categories like "sidewall-graspable"), so the classification itself determines how to grasp the object. Also, as it has been showed by [8], the appropriate usage of grasps is not necessarily related just to the object's shape. For example, the way to grasp a hammer is not the most natural or most stable for this object, but it is the best for the purpose a hammer is used.

According to [9], a grasp action involves two main functions well separated: the approach component (involving the arm muscles) and the grasp component (involving the hand muscles). Although it has been showed that these systems

are closely correlated, people focused mainly on one of the two subsystems. There are systems performing imitation of the arm [10] or, more generally, the upper-body [11]. The arm/upper-body imitation does not experience the self occlusion to the same extent as the hand does.

Our previous work [4] considered an HMM framework for recognition of grasping sequences using magnetic trackers and evaluated both the fingertip and the posture mapping. There are approaches that imitate the whole hand posture [12] or perform a simple mapping to a gripper based on two fingertips [13]. Mapping grasps to more complex hands is usually much more complicated. In [14] the concept of "virtual finger" is introduced: one or more real fingers acting in unison. Kang and Ikeuchi [15] use this concept as an intermediate step in their mapping procedure. Our approach integrates vision based mapping and the notion of virtual fingers for mapping human grasps to two robot hands: one of them resembles human kinematics (Karlsruhe hand) and one of them does not (Barrett hand). Thus, the robot here does not explore a range of approach vectors, but instead directly imitates the human approach vector, encoded in the hand position and orientation relative to the object.

III. Vision Based Grasp Classification

The details of our methodology for visual recognition of grasps can be found in our previous work, [5]. We only summarize the most important aspects. The input to the grasp classification method is a single monocular image in which we first segment the hand. The classification method is non-parametric; grasp classification and hand orientation regression is formulated as a problem of finding the hand poses most similar to image example $\mathbf{H}$ in a large database. Each database sample $\mathbf{H}_{i,j}^{\text{synth}}$, where i denotes grasp type and j denotes sample number, has associated with it a class label $y_{i,j} = i$ and a hand-vs-camera orientation $o_{i,j} = [\phi_j, \theta_j, \psi_j]$, i.e. the Euler angles from the camera coordinate system to a hand-centered coordinate system. To find the grasp class $\hat{y}$ and orientation $\hat{o}$ of an unknown grasp view $\mathbf{H}$ acquired by the robot, a distance-weighted k-nearest neighbor (kNN) classification/regression procedure is used. First the set of k nearest neighbors to $\mathbf{H}$ in terms of Euclidean distance between gradient orientation histograms obtained from the grasp images are retrieved from the database. From the found approximate nearest neighbors, the estimated class of $\mathbf{H}$ is found as a distance-weighted selection of the most common class label among the k nearest neighbors, and the estimated orientation as a distance-weighted mean of the orientations of those samples among the k nearest neighbors for which $y_{i,j} = \hat{y}$.

IV. Example-Based Mapping of Grasp to Robot

The estimated grasp class as well as hand and object orientation and position are used to instantiate a robot grasp strategy. The human-to-robot grasp mapping scheme is defined depending on the type of robot hand used. The Barrett hand is a three fingered, 4DOF robotic hand with an embodiment substantially different to a human hand. Karlsruhe hand is a five fingered, 8DOF robotic hand with an embodiment similar to a human hand. The preshapes used for the hands are shown in Fig. 1. There are three preshapes for the Barrett hand:

- Barrett Wrap: used for grasps with a preshape with large aperture, like Large and Small Diameter grasps;
- Barrett Precision Grasp: for small aperture preshapes like the Pinch grasp and the Abducted Thumb (executed as a pinch grasp due to the hand kinematic constraints);
- Barrett Precision Disc: for circular objects.

There are four preshapes for the Karlsruhe hand:

- Karlsruhe Power preshape is applied for grasps with four parallel fingers and thumb opposed to them, like Large Diameter, Pinch and Abducted Thumb.
- Karlsruhe Wrap is applied for the Small Diameter where the thumb is not opposed to the rest of the fingers.
- There are two preshapes for round objects, Karlsruhe Sphere (for Power Sphere) and Karlsruhe Disc (for Precision Disc); the differences are in the pose of the thumb (more opposed in the Disc) and in how straight are the rest of the fingers (more bent in Power Sphere).

The grasp mapping is performed as shown in Algorithm 1. The hand orientation estimate $o_{h\to c}$, along with the hand position estimate $p_{h\to c}$ and the estimated position and orientation $o_{o\to c}, p_{o\to c}$ of the grasped object all relative to the camera are used to derive the estimated position and orientation of the human hand relative to the object $o_{h\to o}, p_{h\to o}$. The hand orientation relative to the table plane o_h is extracted from $o_{h\to c}$ and the orientation of the camera o_c, obtained through the robotic head kinematics. The estimation of object position and orientation is assumed perfect; this part of the system is not implemented, instead the ground truth is given in the simulations.

The system first decides which preshape to use based on the recognized grasp. Then, the approach vector is chosen. Two different ways of approaching the object are used, based on the orientation of the human hand; if the palm orientation o_h is similar to the one in Fig. 1(e) the object is approached from the side, otherwise it is approached from the top. Based on the estimated type of grasp, the system differentiates between volar and non-volar grasp, [15], i.e., whether there should be a contact between the palm and object or not. The original volar grasps are the Large Diameter, Small Diameter, Abducted Thumb and Power Sphere grasps, see Fig. 1. However, the limitations of the hands embodiments make impossible to use the palm in the Abducted Thumb and Power Sphere grasps. In a human Abducted Thumb grasp the palm adapts its shape to the object, and the abduction/adduction degrees of freedom of the fingers are used; the robotic hands studied here lack those degrees of freedom, so the Abducted Thumb is mapped to a Pinch Grasp. In the case of the Power Sphere, the robotic hands cannot apply a volar grasp due to the larger length of robotic fingers.

The volar grasping is performed in the following order:

1) The robot adopts the hand orientation and preshape

Data: Human Grasp G_h, $p_{h\rightarrow o}$,$o_{h\rightarrow o}$, o_h
/* Robot Hand $H \in \{Barrett, Karlsruhe\}$ */
/* Robotic Grasp G_r */
/* Approach Vector a */
/* Distance palm-fingertip δ */
if $H = Barrett$ **then**
  if $G_h \in \{LargeDiameter, SmallDiameter\}$ **then**
    $G_r = BarrettWrap$;
  else if $G_h \in \{Pinch, Abducted\}$ **then**
    $G_r = BarrettPrecision$;
  else if $G_h \in \{PowerSpherical, PrecisionDisc\}$ **then**
    $G_r = BarrettPrecisionDisc$;
else if $H = Karlsruhe$ **then**
  if $G_h \in \{LargeDiameter, Pinch, Abducted\}$ **then**
    $G_r = KarlsruhePower$;
  else if $G_h = SmallDiameter$ **then**
    $G_r = KarlsruheWrap$;
  else if $G_h = PowerSphere$ **then**
    $G_r = KarlsruhePowerSphere$;
  else if $G_h = PrecisionDisc$ **then**
    $G_r = KarlsruhePrecisionDisc$;
if o_h = o_{side} **then** ; /* see Fig. 1(e) */
  $a = a_s$; /* approach from side */
else
  $a = a_v$; /* approach from top */
/* execute */
Set hand to preshape G_r;
Set hand to orientation $o_{h\rightarrow o}$;
if $G_h \in \{LargeDiameter, SmallDiameter\}$ **then** ;
/* if volar */
  Approach following a towards $p_{h\rightarrow o}$ until contact;
  while *contact* p_c *out of palm* **do**
    Retreat;
    $p_{h\rightarrow o} = \alpha p_c + (1-\alpha)p_{h\rightarrow o}$;
    Approach following a towards $p_{h\rightarrow o}$;
    $\alpha = \alpha^2$;
else ; /* if non-volar */
  Approach following a towards $p_{h\rightarrow o} - \delta$;
Grasp;

Algorithm 1: Pseudo-code for grasp mapping.

corresponding to the estimated human grasp.

2) The robot hand approaches the object centroid until it detects contact on the palm sensor. After that, it closes the hand.

When the robot detects that the first contact did not occurred in the palm, the trajectory is replanned. The new goal position for the hand is a weighted average between the detected contact p_c and the original goal position $p_{h\rightarrow o}$, as explained in Algorithm1. The non-volar grasps, which have no contact between the palm and the object, are originally the Pinch and Precision Disc grasps (see Fig. 1). Since there is no contact between the tactile sensor in the palm and the object, in our system the grasp is performed without any feedback. For this reason, the non-volar grasps depend heavily on the precision of the object position and orientation estimation. The difference between non-volar and volar strategies is the absence of the loop where the contact location is checked (see Algorithm 1).

V. Experimental Results

We first evaluate our approach in the GraspIt! simulator and then demonstrate it in a real robotic scenario.

A. Simulated grasping with GraspIt

As stated, evaluating the performance of a grasp imitation system is not trivial. It cannot be based on grasp stability, and comparison between joint angles in robotic hands and human hand is not possible because of the differences in the embodiment. We have decided to compare the grasps using the concept of virtual fingers ([14]), computed based on the equations stated in [16]. As cited in Section II, a virtual finger is a group of real fingers (including the palm) that act in unison. So, in theory, the average position and orientation of the virtual finger contacts in the imitated grasp should be similar to the ones in the original grasps. However, as it will be discussed later, that is not always the case.

The virtual finger configuration tries to minimize two factors: the number of virtual fingers N, in order to achieve a compact representation, and the heterogeneity of the real fingers R_i conforming each virtual finger V_k, described as a cohesive index for virtual finger k, C_{V_k}. The cohesive index of each virtual finger is computed based on the degree of force coupling (cosine of the angle between the forces) between each two forces f_i, f_j applied with any of the fingers within a virtual finger:

$$\begin{aligned} D_c(i,j) &= \frac{f_i \cdot f_j}{|f_i| \cdot |f_j|}, \;\; m_{ij} = \frac{1 + D_c(i,j)}{2} \\ C_{V_k} &= \prod_{\substack{i \in R_i, j \in R_j \\ R_{i,j} \in V_k}} m_{ij}^{\xi}, \;\; \xi = \binom{F(V_k)}{2}^{-1} \end{aligned}$$

where $F(V_k)$ is the number of fingers within V_k. For example, if all forces within a virtual finger k are parallel, $C_{V_k} = 1$. If any two forces belonging to the virtual finger are perpendicular, $C_{V_k} = 0$. So, in order to find the best configuration of virtual fingers, we maximize the cohesive indexes C_{V_k} trying to keep the number of virtual fingers small:

$$\begin{aligned} \text{Maximize} \quad & C_{\mathit{eff}} = (\tfrac{1}{N!} \textstyle\prod_{i=1}^{N} C_{V,i})^{\frac{1}{N}} \\ \text{Subject to} \quad & N \in 1,2,3,4,5,6, \;\; \textstyle\bigcup_{i=1}^{N} V_i = R \\ & V_i \cap V_j = \varnothing, i \neq j, 1 \leq i,j \leq N \end{aligned}$$

So for every possible combination of real fingers R_i assignment to virtual fingers V_k, the coefficient C_{eff} is

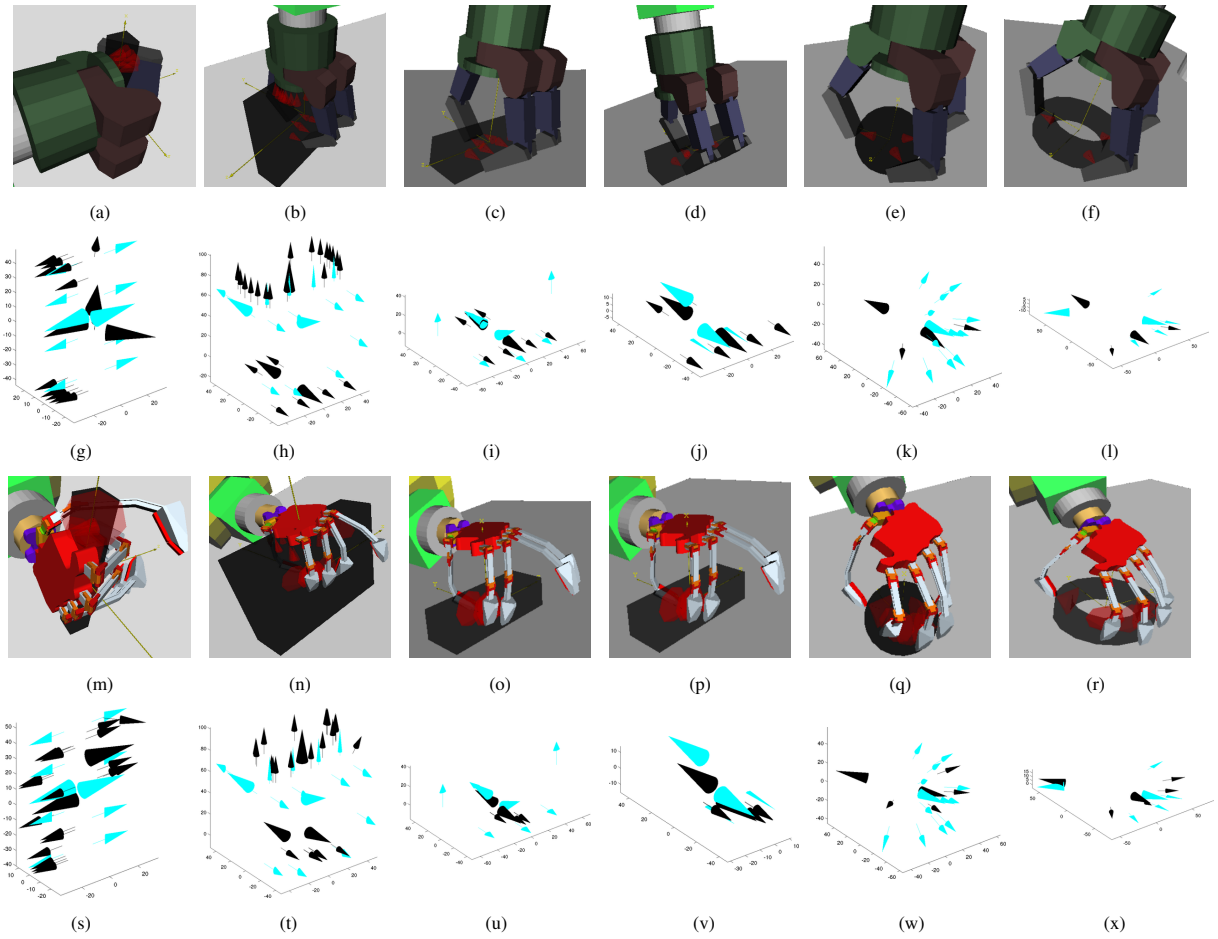

Fig. 2. First and third row shows the grasp execution in absence of errors for Barrett and Karlsruhe Hand. Second and forth row show a comparison between the contacts for Barrett(black)-Human(blue) hands and Karlsruhe(black)-Human(blue) hands. The big arrows show the average pose of the virtual fingers,$\mathcal{P}_{V_k}$

computed. The assignment with the highest C_{eff} is selected as the virtual finger representation of the grasp. The position and orientation of contacts is automatically extracted from the robotic simulator for the robot grasps, and it was tagged manually from images for the human grasp.

So far, all the real fingers R_i have been assigned to a virtual finger V_k. In order to compare the configuration of different hands we will define the contact pose (6d, including position and orientation) $\mathcal{P}_{V_k}$ of a virtual V_k as the average of the contact poses $\mathcal{P}_i$ within this virtual finger:

$$\mathcal{P}_{V_k} = \frac{1}{T(V_k)} \sum_{\substack{i \in R_i \\ R_i \in V_k}} \mathcal{P}_i$$

$$T(V_k) \equiv \text{number of contacts within } V_k$$

In the first experiment, perfect object pose estimation and perfect hand pose recognition is assumed. Fig. 2 represents the grasps and the contact comparison between the robotic hands (black) and the human hand (blue). The big arrows show $\mathcal{P}_{V_k}$, and the small arrows show all $\mathcal{P}_i$. It can be seen in this figure that the pose of the virtual fingers and even the number of them does not coincide always. For example, the Barrett hand has three virtual fingers for the Small Diameter grasp, while Human and Karlsruhe hands have two (Fig. 2a,g,m,s). The reason for this mismatch is that Barrett fingers are longer than human and Karlsruhe fingers, so the object is touched by the last phalanx in the edges instead of the face. Another significant difference in the number of virtual fingers appears in the Power Sphere grasp (Fig. 2c,k,q,w). The human grasp has just one virtual finger, while the robotic hands have two. For the human, placing the thumb opposed to the rest of the fingers is uncomfortable. The big contact surface and therefore big friction between the hand and the ball allows him to place the fingers in a relatively unstable way. However, the contact surface between the robotic hands and the object is much smaller, so the thumb should be placed in opposition to the rest of the fingers. In contrast, in the Precision Disc grasp (Fig. 2f,r,l,x) the human needs to place the thumb in opposition to the rest of the fingertips due to the lower friction between the hand and the object. It is also interesting to reason about the results for the Large Diameter grasp (Fig. 2b,h,n,t). Apparently there is a big difference between the average virtual finger position, but the actual contacts look similar. The reason is that the human fingers have contacts in both the proximal and distal phalanges, while

the robot achieves a contact just in the distal "phalanges". Finally, it should be pointed that despite the impossibility of imitating properly the Abducted Thumb grasp, the position of the virtual fingers is quite accurate, with a deviation in orientation in one of them due to the two contacts from the human palm and index fingertip in the top of the object, inexistent in the robotic grasps.

We present a more concise view of the experiments in Fig. 3: it represents the average error in virtual fingers position and orientation (position and orientation component of $\mathcal{P}_{V_k}$) between the robotic hands (where Barrett error is represented in black and Karlsruhe in white) and the human hand. However, it should be noted that this measure is a lower bound of the error: in cases where the number of virtual fingers is different, this represents the average distance between the best matching virtual fingers. Sometimes this mismatch is known and natural (like the different number of virtual fingers in the Power Sphere grasp), but sometimes this means that one finger failed to touch the object. This happens mainly in the experiments with position error with the Karlsruhe hand, and will be mentioned later. For the case of perfect data (Fig. 3a for orientation, Fig. 3e for position) we can infer a number of conclusions: Karlsruhe hand performance in terms of orientation is better than the performance of Barrett hand; the performance in terms of position of the virtual fingers is similar; the biggest errors appear in the Large Diameter grasp, for the reasons stated before.

The next experiment tests the robustness with respect to the object position errors. We introduced an error of $\rho = 5cm$ in 6 different directions:

$$\begin{aligned} \widehat{p_o} &= p_o + \rho[\cos\beta, \sin\beta, 0] \\ \beta &= \{0, \frac{\pi}{3}, 2\frac{\pi}{3}, 3\frac{\pi}{3}, 4\frac{\pi}{3}, 5\frac{\pi}{3}\} \end{aligned}$$

The difficulty of the problem should be noted: the size of the objects in their biggest axis is around 10cm, so the error is significant compared to their size. Another factor is the lack of any visual feedback. This experiment show us principally the importance of the feedback (tactile feedback in our case) in the presence of errors. The only grasps where tactile feedback was used are the Large and Small Diameter grasps. The reason for that is because those grasps are the only ones where we expect a first contact with the palm: this means that a first contact detected in any other finger suppose an error in the object pose detection that should be corrected. It can be seen that, in the presence of object pose errors, the error increases much more in the grasps without corrective movements (grasps 3,4,5 and 6) than in the ones with corrective movements (grasps 1 and 2). Another thing that can be inferred in Fig. 3b,f is that the Karlsruhe hand is more sensitive to the errors that the Barrett hand; the error increases more for Karlsruhe hand in presence of errors than for the Barrett hand. Actually the error for the Karlsruhe hand in non-corrected grasps is higher than the one showed, because the thumb usually fail to touch the object, and therefore the thumb virtual finger not compared. There are principally two reason for this worse robustness to errors: first, the shorter length of the fingers; second, the palm configuration in the Karlsruhe hand. The shorter length of Karlsruhe fingers affect the non-volar grasps, as we can see in Fig. 4a,b. The small distance between the base of the thumb and the base of the rest of the fingers affects the volar grasps, that usually collide with the finger bases before touching the palm. However, this is mostly solved by the corrective movements.

(a) (b)

Fig. 4. Example of performance of a grasp without corrective movements

B. Real grasping with a KUKA arm

An image of the human grasp is captured and passed to the grasp recognition module [5], which returns the type of grasp and the position and orientation of the hand. Using the object pose, the grasp policy is selected and executed. The scenario, illumination and subject is different to the experiments in [5], but we get similar results in the classification. Large diameter, small diameter and abducted thumb are correctly classified most of the time, while pinch grasp, power sphere and precision disc grasp are sometimes confused with the power grasp. In terms of orientation, the typical error is around 15 degrees, which is acceptable in the execution of the grasp. The object position is given manually, with an error of ± 3 cm. The position error did not inflict on the grasp execution, except when performing Precision Disc grasp with a ball, which rolled when the hand was not centered over the ball. Fig. 5 shows the robot being shown four different grasps (Large Diameter, Abducted Thumb, Pinch and Precision Disc, respectively), mapping them and performing the corresponding grasp (Barrett Wrap, Barrett Precision, Barrett Precision and Barrett Precision Disc, respectively).

VI. Conclusions

We have presented a human-to-robot grasp mapping system based on a single image. We have proposed an evaluation metrics for assessing the quality of mapping. The approach was demonstrated both in simulation and in a real robot setup. The system presented here can be improved in several ways. The database of hand poses should include more objects of different sizes, and more advanced non-parametric regression methods could be employed for estimating grasp type and hand pose. Moreover, the addition of visual servoing would improve considerably the performance in grasps without tactile feedback (volar grasps) or grasps when the

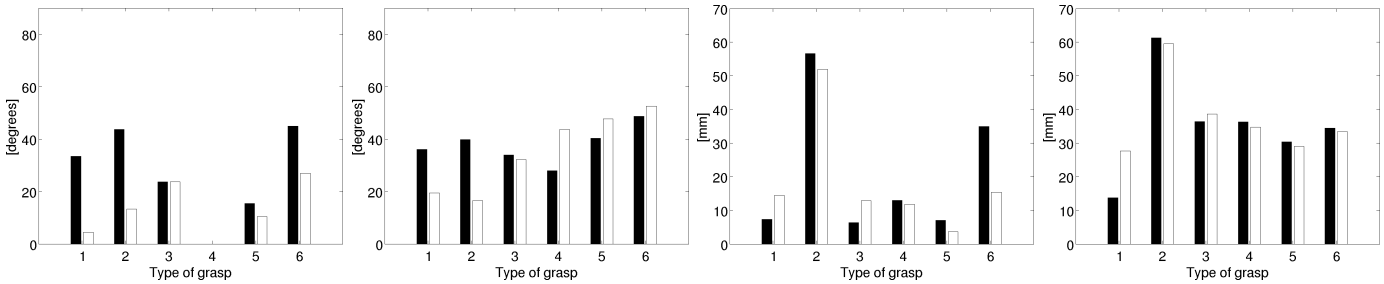

(a) Orientation error: perfect pose (b) Orientation error: 5cm pose error (c) Position error: perfect pose (d) Position error: 5cm pose error

Fig. 3. Error in position (mm) and orientation (degrees) for each of the six grasps tested (Small Diameter, Large Diameter, Abducted Thumb, Pinch, Power Sphere, Precision Disc). Each column represents experiments with no error and error of 5cm.

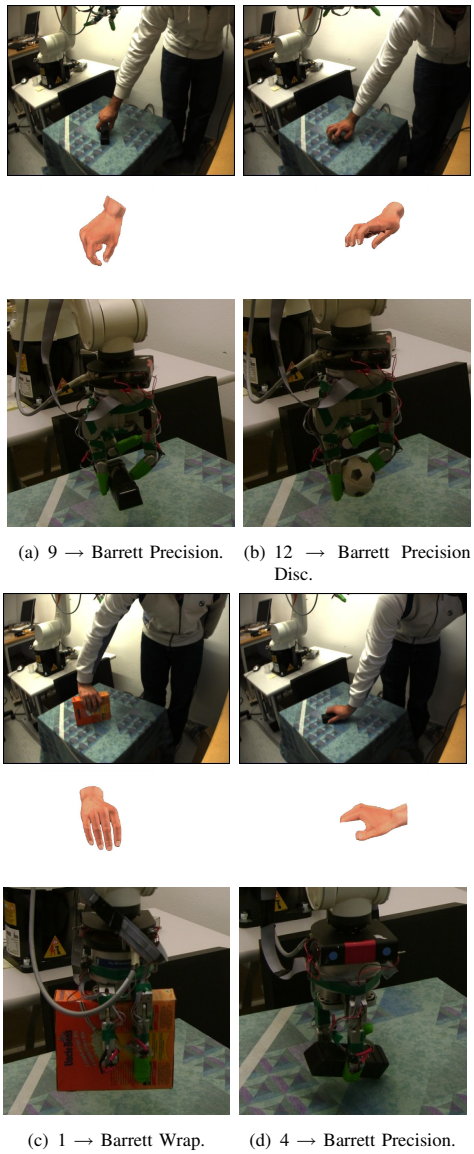

(a) 9 → Barrett Precision. (b) 12 → Barrett Precision Disc.

(c) 1 → Barrett Wrap. (d) 4 → Barrett Precision.

Fig. 5. Execution of grasps in a real robot environment: original images, nearest neighbors in the database and robot execution.

tactile feedback fail due to the limitations of the hand sensors. Finally, we will investigate the performance using a humanoid robot hand.

REFERENCES

[1] Y. Kuniyoshi, M. Inaba, and H. Inoue, "Learning by watching," *IEEE Transactions on Robotics and Automation*, vol. 10, no. 6, pp. 799–822, 1994.

[2] S. Schaal, "Is imitation learning the route to humanoid robots?" *Trends in Cognitive Sciences*, vol. 3, no. 6, pp. 233–242, 1999.

[3] K. Ogawara, S. Iba, H. Kimura, and K. Ikeuchi, "Recognition of human task by attention point analysis," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2000, pp. 2121–2126.

[4] S. Ekvall and D. Kragić, "Grasp recognition for programming by demonstration tasks," in *IEEE International Conference on Robotics and Automation*, 2005, pp. 748–753.

[5] H. Kjellström, J. Romero, and D. Kragić, "Visual recognition of grasps for human-to-robot mapping," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.

[6] M. Cutkosky, "On grasp choice, grasp models and the design of hands for manufacturing tasks," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 3, pp. 269–279, 1989.

[7] M. Stark, P. Lies, M. Zillich, J. Wyatt, and B. Schiele, "Functional object class detection based on learned affordance cues," in *Computer Vision Systems*, 2008.

[8] A. M. Borghi, "Object concepts and action," in *The Grounding of Cognition: The role of perception and action in memory, language, and thinking*, D. Pecher and R.Zwaan, Eds. Cambridge University Press, 2005, part 2, pp. 8–34.

[9] M. Jeannerod, "Intersegmental coordination during reaching at natural visual objects," *Attention & Performance*, vol. IX, 1981.

[10] A. Billard, Y. Epars, S. Schaal, and G. Cheng, "Discovering imitation strategies through categorization of multi-dimensional data," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, 2003, pp. 2398–2403.

[11] P. Azad, A. Ude, T. Asfour, and R. Dillmann, "Stereo-based markerless human motion capture for humanoid robot systems," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 3951–3956.

[12] J. Rehg and T. Kanade, "Visual tracking of high dof articulated structures: An application to human hand tracking," in *European Conference on Computer Vision*, vol. 2, 1994, pp. 35–46.

[13] J. Triesch, J. Wieghardt, E. Mael, and C. Malsburg, "Towards imitation learning of grasping movements by an autonomous robot," in *International Gesture Workshop*, 1999, pp. 73–84.

[14] M. A. Arbib, T. Iberall, and D. M. Lyons, "Coordinated control programs for movements of the hand," in *Hand function and the neocortex. Experimental Brain Research Supplemental 10*, A. W. Goodwin and I. Darian-Smith, Eds., 1985.

[15] S. B. Kang and K. Ikeuchi, "Toward automatic robot instruction from perception: Mapping human grasps to manipulator grasps," *IEEE Transactions on Robotics and Automation*, vol. 13, no. 1, pp. 81–95, 1997.

[16] S. B. Kang and K. Ikeuchi, "A framework for recognizing grasps," in *CMU Robotics Institute*, 1991. [Online]. Available: ftp://reports.adm.cs.cmu.edu/usr/anon/robotics/CMU-RI-TR-91-24.ps.Z