

# A Compact Representation of Human Single-Object Grasping

Steffen Puhlmann

Fabian Heinemann

Oliver Brock

Marianne Maertens\*

**Abstract**—Observations of human grasping reveal that the exploitation of environmental constraints is a key structural aspect for the robustness and versatility of human grasping behavior. We analyze 3,400 human grasping trials with 17 subjects grasping 25 objects to show that viewing environmental constraints as the central structural aspect of human grasping yields surprisingly simple representations of human grasping behavior. We present hypothesis-driven experiments that emphasize the centrality of environmental constraints in human grasping and extract from data a simple “grasping plan” that is a generative model for all of the human grasping trials we observed. This grasping plan can in principle be transferred to a robot system in an attempt to leverage environmental constraints to improve the performance of robotic grasping.

## I. INTRODUCTION

Human grasping is far superior to robot grasping. We investigate the principles of human grasping that lead to this superior performance and work towards transferring these principles to robotic systems. The key hypothesis of our research is that grasping performance in humans crucially depends on the purposeful exploitation of contact with the environment. The environment provides *environmental constraints* (ECs), whose exploitation through contact during grasping relaxes the requirements on accurate control, simplifies perception, and facilitates grasp planning. Our results demonstrate that a successful replication of human-level grasping performance in robotic systems must carefully consider the exploitation of environmental constraints.

We analyze 3,400 single-object grasping trials to show that sequences of EC exploitations yield a compact and expressive representation of human grasping behavior. These sequences consist of action primitives connected by sensory events, indicating the termination of a primitive and the invocation of the next. In prior work, we had shown that only six action primitives are necessary and sufficient to describe the observed behavior (*reach, close, slide, edge-grasp, flip, and fail*; see Fig. 5 in reference [1]). In this paper, we characterize the sequences of primitives that occur in human grasping. We find that simple transition graphs capture the observed sequencing. The simplicity of the resulting representation is an indication that ECs capture an important aspect of human grasping behavior.

All authors, except for Marianne Maertens, are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany.

\*Modelling of Cognitive Processes Group, Technische Universität Berlin

We gratefully acknowledge the funding provided by the Alexander von Humboldt foundation and the Federal Ministry of Education and Research (BMBF), by the European Commission (EC, SOMA, H2020-ICT-645599), the German Research Foundation (DFG, Exploration Challenge, BR 2248/3-1), and by an Emmy-Noether research grant of the German Research Foundation to Marianne Maertens (DFG, MA5127/1-1).

The evaluation of human grasping trials also demonstrates the importance of environmental constraints for grasping performance. Our experiments reveal that humans take advantage of ECs to counteract uncertainty. We also show that there is significant deterioration in grasping performance when humans are instructed to avoid the use of environmental constraints. Both of these results indicate that human grasping capabilities critically depend on the use of ECs.

Taken together, the effectiveness of ECs in describing human behavior and the importance of ECs for human grasping performance indicate that the exploitation of environmental constraints is a central consideration for achieving human-level grasping performance in robotic systems.

Towards the goal of replicating human grasping principles in robotic systems, we identify the sensory requirements for detecting transitions between consecutive primitives. Our analysis shows that the observed transition conditions can be realized on existing robotic hardware. Together with the fact that the action primitives are also transferable to robots [1], our results provide evidence for the feasibility of transferring principles of human EC exploitation to robotic systems.

## II. RELATED WORK

Human centered taxonomies in robotics [2], [3], [4] define static hand postures, i.e. they specify postures of the hand that are categorized according to grasp strength or object shape. Such static descriptions are interesting for the analysis of contact configurations but neglect the dynamic nature of attaining that configuration. In contrast, we believe that dynamic aspects resulting from the interactions between hand, object, and environment are essential for devising successful and robust approaches to robot grasping. In particular, we study human EC-exploiting pre-grasp manipulations [5], [6] and grasping motions, because we assume that they contribute to robustness and versatility of human grasping. We therefore propose that an EC-based taxonomy [1] is the appropriate framework to analyze human grasping.

The extraction and transfer of human grasping knowledge to robotic grasping has previously been demonstrated by having humans physically guide the robot into a grasping pose [7]. In our work, we directly study grasping based on the human hand.

Worgötter [8] suggested to provide structure to an agent by considering how manipulation actions are sequenced in time. They were surprised to find a rather restricted manipulation space and a simple tree-like structure for most manipulations. Mandery [9] analyzed the sequencing of support poses for whole body movements. They showed that the use of support poses increased robustness especially in the face

of uncertainty. Inspired by the above results we scrutinize the sequencing of action primitives when humans exploit ECs during grasping. Our results confirm that, given an appropriate structural emphasis, seemingly diverse behaviors derive from a simple representational basis.

Previous experiments showed that sensory uncertainty leads to an increased use of environmental constraints (ECs) [10], [11]. These findings strongly suggest that ECs are important, but the increased use of support contacts could be a by-product of some other strategy, instead of being a strategy in its own right. To demonstrate that humans willfully exploit ECs to achieve robust performance, we tested human grasping when EC exploitation was not possible. Inspired by Kazemi [10], we included a condition in which subjects were not allowed to touch the support surface because it was supposedly hot. This lead to radical changes in strategy and increased failures, supporting our assumption that the exploitation of ECs is not a by-product but instead an important contributor to human grasping capabilities.

### III. DATA COLLECTION

We study human grasping in a single-object table-top scenario (Fig. 1). In each trial, the participant grasps one out of 25 different objects (Fig. 2) placed in front of her and lifts the object in accordance with the experimental protocol (detailed below). At the start of a trial, participants' hands rested at the starting position on the table (Fig. 1). Following an audio signal, the participant initiates the grasp using only the right hand. Data from the trial is recorded starting with the audio signal and ending with the lifting of the object.

Seventeen right-handed subjects (seven female, age range from 23 to 35 years) participated in the experiment. Subjects had no prior knowledge of the purpose of the experiment, and participated in a single experimental session, lasting about two hours. Participants gave informed consent prior to the experiments and the experimental protocol was approved by the Institutional Review Board of the University. Participants received a financial compensation of 8€ per hour.

We used 25 different objects grouped into one of six categories (Fig. 2). The objects in each of the categories were chosen so as to elicit different grasping actions based on observations in a previous experiment [1]. The categories are named based on the targeted grasping behavior, including the category 'new' for strategies that have not been observed in previous experiments:

- **flip:** button, french chalk, key, shell
- **edge grasp:** credit card, CD, comb, game card
- **closing:** salt shaker, tape, toy, chestnut, matchbox
- **pinch:** screw, match, cigarette, rubber band
- **rotation:** marker, screw driver, shashlik, glasses
- **new:** coffee mug, plate, book, bowl

Subjects grasped under different experimental conditions. Analogous to our previous study [1], we varied the amount of sensory uncertainty. In addition to the normal vision condition, we included an impaired vision condition in which subjects wore frosted-glass goggles (Fig. 1). We also manipulated whether or not the participants would use the

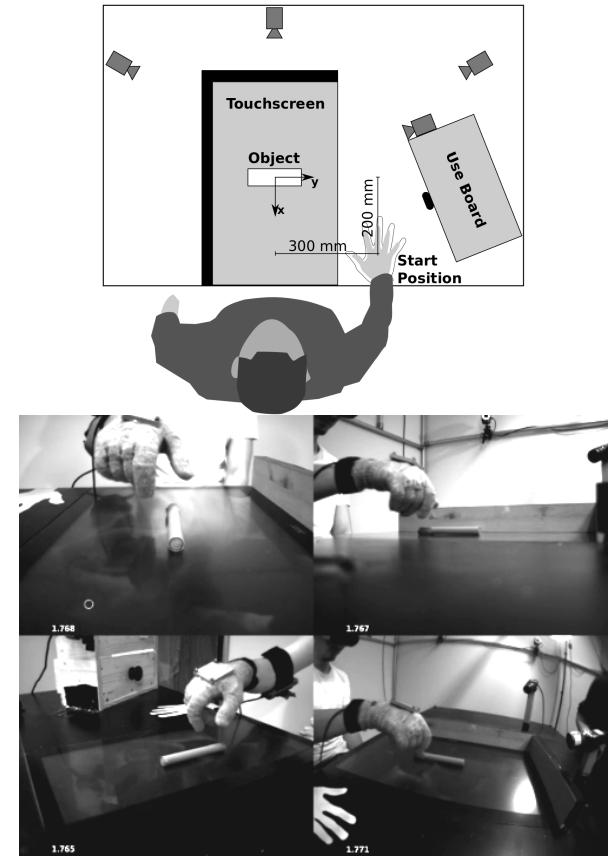


Fig. 1. Experimental setup for the grasping experiments: schematic diagram (top) and participant grasping the marker with impaired vision in the "hot surface" experimental condition (bottom, four viewpoints).

support surface by introducing a hot-surface condition [10]. We told the subjects to imagine that the table is extremely hot and to avoid contact with it. To augment the instructions, we presented an image of burning charcoal on the touchscreen and played a loud noise upon contact. The visual conditions (normal vs. impaired) and the surface conditions (normal vs. hot) were factorially combined. In these four experimental conditions, observers were instructed to grasp, lift, and hold the object. Each object was grasped twice in each of the five experimental conditions, resulting in a total number of 200 trials per participant.

### IV. DATA ANALYSIS

#### A. Segmentation and annotation

We manually segmented the video-recordings using the action primitives defined in our previous EC-based taxonomy [1]. We start by using human labelling instead of data driven approaches [12], [13], [14] because, following Mason [15], we assume that humans are quite proficient at recognizing actions in other people's motion.

We used a customized graphical user interface to facilitate annotation (similar to the one described in [15]). It displays the videos from all four cameras in the center of the GUI window (Fig. 1). The labeller can start and stop videos manually, and can navigate through the the trial frame by frame. To

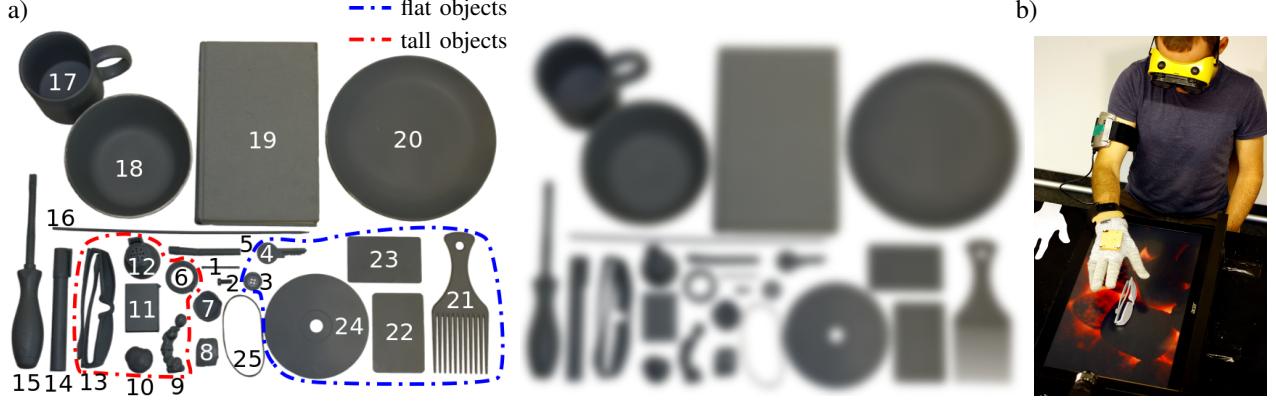


Fig. 2. a) 25 Objects used in the experiment: (1) match, (2) screw, (3) button, (4) key, (5) cigarette, (6) tape, (7) shell, (8) french chalk, (9) toy, (10) chestnut, (11) matchbox, (12) salt shaker, (13) glasses, (14) marker, (15) screwdriver, (16) shashlik, (17) coffee mug, (18) bowl, (19) book, (20) plate, (21) comb, (22) game card, (23) credit card, (24) CD, and (25) rubber band viewed with normal (left) and impaired vision (middle). b) Participant grasping the glasses with impaired vision on “hot” surface (right)

annotate grasping actions the start and stop timestamps of the corresponding video segment are recorded and saved together with its given annotation. Two people, one experienced (author F. H.) and one inexperienced, annotated the videos.

#### B. Adequacy of the EC-based taxonomy

In a previous paper [1] we argued for the adequacy of the proposed taxonomy based on the criteria of reliability and validity. We will provide further evidence for the validity of an EC-based representation below. Reliability is assessed as the amount of agreement between the annotations of the two independent labellers. We computed the Levenshtein distances [16] between their annotations. In 75% of the trials the labelled action primitives had a Levenshtein distance of zero, i.e. exactly the same labelling. In 9% of the trials, the action primitives had a Levenshtein distance of one. These values are comparable to those in the previous paper where the respective values were 74% and 15% and two experts labelled the data.

According to the EC-based taxonomy, each grasping trial can be represented by a combination of five different action primitives: *approach* (a), *close* (c), *manipulation* (m), *flip* (f), *edge grasp* (e). They result either in a successful grasp (\*) or they fail (x). Considering the enormous number of different low-level hand configurations during grasping, it is a remarkable feature of the proposed taxonomy that a limited set of primitives captures the observed grasping behaviors.

#### C. Transitions between action primitives

To compute the transition probabilities between action primitives we counted the number of occurrences of one action primitive following another and represented them in a cross-table. We normalized the occurrence of each transition by the total number of transitions for a particular action primitive. From the transition probabilities we created Markov chains, i.e. directed graphs with nodes for each action primitives that are spatially arranged according to their sequence of occurrence, and that are connected by lines of varying thickness to represent their transition probabilities.

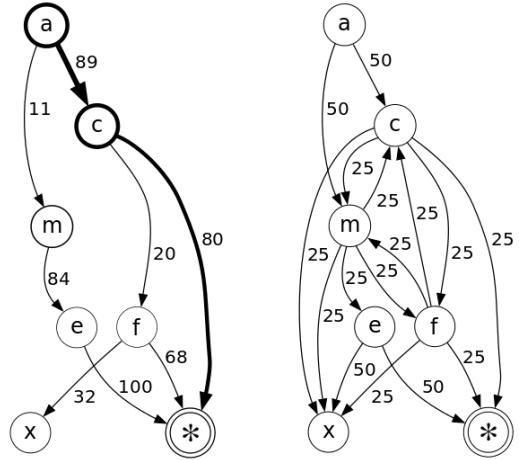


Fig. 3. Markov chains representing transition probabilities between action primitives. Left: Markov Chain describing the sequence of action primitives in the condition with full visual information and interaction with the environment. Including all objects this data set contains 850 trials (17 subjects, 25 objects, 2 repeats). Right: Markov chain describing all possible transitions that would be consistent with our taxonomy. Transition probabilities are uniformly distributed between successor states. We only show transitions with a probability of occurrence of more than 5%.

First, we examine human grasps under standard conditions with unrestricted visual information and allowing the interaction with the environment. The model extracted for all objects (850 trials) is shown in Figure 3 as a Markov chain.

We would like to demonstrate that the obtained Markov model indeed captures the structure inherent to human grasping. To support this, we compare the probabilities of the most common grasping paths in the data-derived model with the probabilities of the same paths in a quasi-random Markov chain (Fig. 3, right). The quasi-random Markov chain only contains those transitions consistent with our taxonomy (for example, a transition from approach (a) to flip (f) is not possible, because the fingers need to close before the object can be flipped). The transition probabilities are uniformly distributed for each node.

There is one major route of action in the human model

that is pursued in 71% of all trials. This path is composed of an approach phase (first action in all trials), followed by a closing action (89% of all trials) which results in a successful grasp (80% of the cases). In the “random” Markov chain, the analogous probability for the path preferentially chosen by our subjects is about 12% (combining the transition probabilities from (a) to (c) and from (c) to (\*)). We view this as support for our hypothesis that our representation of human grasping behavior captures important structure aspects: in the human model, a single path account for a large number of *successful* trials. In contrast, the same path has a much lower probability in the random (structure-less) Markov chain.

To complement the above analysis, we also evaluate the statistical significance of the difference between the two graphs in Fig. 3. For this, we employ binomial tests. Specifically, we compare the probabilities for the three most frequently observed action sequences in the estimated Markov chain with the corresponding probabilities in the random Markov chain. The most frequently observed action sequences were (approach - closing - \*), (approach - manipulation - edge grasp - \*), and (approach - closing - flip - \*).

We used the generative properties of Markov chains to generate test samples for the components of each of these paths. We performed Bernoulli experiments with the probability of  $p$  for one action sequence and the complementary probability  $1 - p$  for all other possible sequences. We repeated these Bernoulli experiments 850 times corresponding to the number of trials that entered the analysis. With the resulting variability estimate, we performed the binomial test, and it revealed that the probability for the random Markov chain to produce the same outcome as the estimated Markov chain is far below 5% (the threshold usually considered for statistical significance in these tests).

The random transition graph, when compared to the data-derived transition graph is therefore a rather unlikely model of the observed grasping behavior. Again, we can interpret this as an indication that our representation captures structure inherent to successful human grasping trials.

#### D. Transition graphs as a function of the objects

In previous experiments [1], we noticed that taller objects were grasped from the side whereas flat objects were grasped with a flip or an edge grasp. Here we tested the reproducibility of this observation by analyzing the same data as in Figure 3 (left) separately for flat and tall objects. We assigned the following objects to the respective categories. Flat objects are the button, the key, the comb, the game card, the credit card and the CD (red-rimmed objects in Fig. 2), and tall objects are the tape, the toy, the chestnut, the matchbox and the salt shaker (blue-rimmed objects in Fig. 2).

Figure 4 shows the transition probabilities for flat and tall objects and they evidently confirm our observation. All 816 trials in which participants grasped a tall object followed the same plan consisting of approach, closing and successful grasp (Fig. 4, right). For flat objects there were two main paths to a successful grasp consisting either of

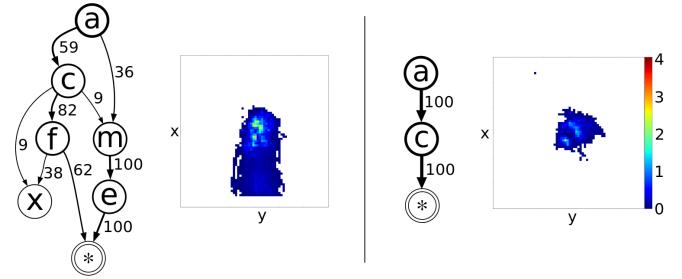


Fig. 4. Transition graphs and contact traces (mean touches per trial) for flat objects (left) and tall objects (right). Each of the two data sets contains 816 trials (6 objects x 2 repeats x 4 conditions x 17 subjects)

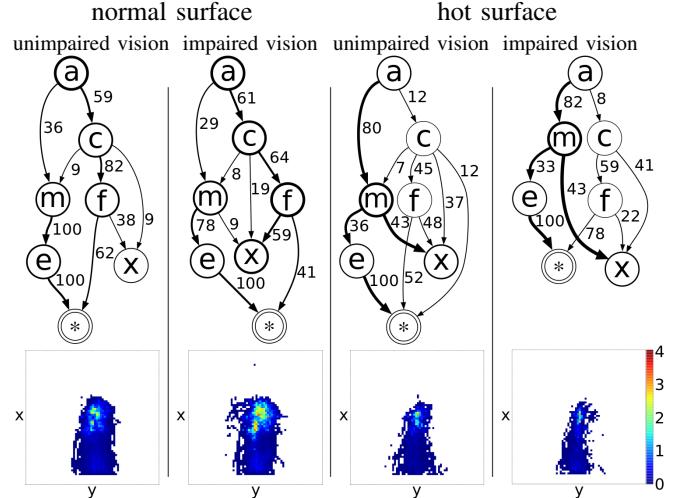


Fig. 5. Transition graphs and contact traces (mean touches per trial) for flat objects in all conditions. Each data set contains 204 trials (6 objects x 2 repeats x 17 subjects).

approach, manipulation and edge grasp (36% of all trials) or of approach, closing and flip (30% of all trials). In 24% of the trials, the grasp action resulted in failure. The different strategies are also evident in the contact traces. They show a higher probability of contact for flat objects and a larger spread towards the edge of the touch screen resulting from ‘sliding’ (manipulation) in combination with ‘edge grasp’. The possibility to exploit the environment thus seems more important when grasping flat objects. Therefore, we continue by analyzing how our experimental conditions affect the grasping strategies for flat objects.

#### E. Transition graphs as a function of the environment

Previous studies reported increased support contact in the presence of visual uncertainty and reduced support contact when objects were placed on a supposedly hot surface [10], [1]. The contact traces for our data (Fig. 5) exhibit the same pattern, providing evidence for the effectiveness of our experimental manipulation. Figure 5 shows the transition graphs for flat objects (same data as in Fig. 4) separated for our experimental conditions. The transition graphs for the normal surface condition are similar to the overall graph (Fig. 4), the full and impaired vision conditions (two graphs on the left of Fig. 5) show little difference. The impaired

vision condition results in a higher number of failures than the normal vision condition (10% vs. 0 for manipulation, 19 vs. 9% for closing, 69 vs. 38% for flip, respectively). Considering the contact traces this seems to imply that most of the visual uncertainty can be compensated by more extensive support contact without much need for a change in strategies.

The contrary was true for the hot surface condition, where different grasping strategies emerge when the surface is not to be touched. While the sequence 'approach, manipulation and edge grasp' remains the main path of action (29% of the trials), the sequence 'approach, closing and flip' is no longer employed (3%). Instead, failures occur in 41% of all trials. Again, there is not much difference between the full and the impaired vision condition. The failure rates are comparable (41 and 39%). All differences following closing are not considered because of the small number of trials (12 and 8% out of 204 trials).

These results show that when subjects are instructed to avoid the use of environmental constraints one of their main strategies for successful grasping becomes obsolete and this leads to a significant deterioration of grasping performance.

#### F. Analysis of Failure

In the last part of the results we explore subjects' recovery strategies to gain insights into the dynamics of grasp control. We created new Markov chains specifically for the failure cases. In addition to showing the probability of failure for a given action primitive, they also show the probability of choosing an alternative successor action subsequent to a failure. Figure 6 (upper row) depicts the error graphs in all conditions for the action primitive 'closing'. Case numbers are shown next to each action primitive. When a failure occurs in the unimpaired vision condition and the normal surface there are three alternative successor actions, 'closing', 'flip' and 'manipulation'. They occur with roughly equal probability of 30%. In the impaired vision condition this ratio shifts towards redoing 'closing' (50%), whereas a 'flip' is performed in only 9% of the cases. In the hot surface conditions a 'flip' is not performed at all after a 'closing' has failed. Instead subjects engage in redoing 'closing' in about 75% of the cases. Figure 6 (middle panel) shows the error graphs for 'manipulation'. In the normal surface condition 'manipulation' is unlikely to result in failure (3 and 9% corresponding to 4 and 15 trials), and if so, the 'manipulation' is repeated in about 75% of the trials. In the hot surface conditions 'manipulation' results in failure in 44% of all trials, and the exclusive recovery strategy is to repeat the 'manipulation' (90% of all trials). Figure 6 (lower panel) shows the error graphs for 'flip'. The recovery strategies for failed 'flips' are rather similar across conditions. In the hot surface conditions 'flips' are generally less likely to occur and with impaired vision a failed 'flip' is never followed by a new flip 'attempt', which seems a viable recovery strategy in all other conditions (around 40%) of the trials. Due to differing and sometimes small trial numbers, we refrained from drawing far-reaching conclusions from these data, but

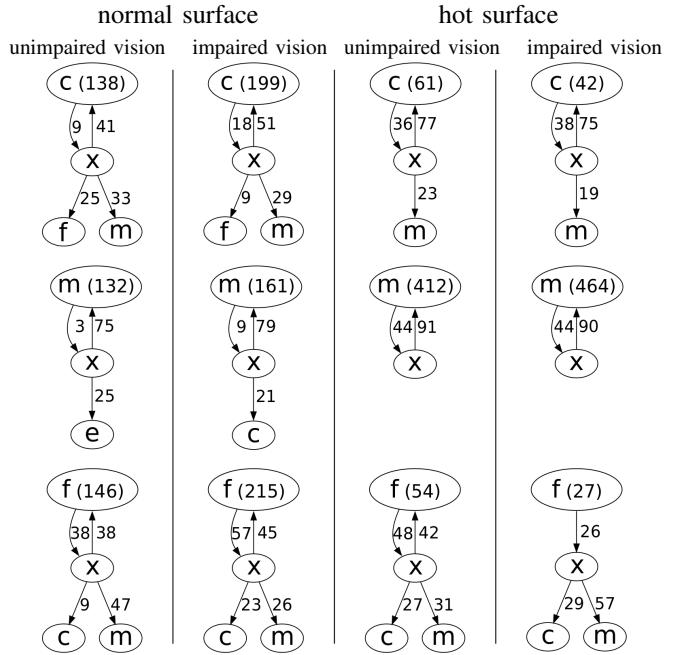


Fig. 6. Error graphs for flat objects in different experimental conditions  
We only show transitions with a probability of more than 5%.

we used them for guidance to identify conditions that might influence the transitions between action primitives.

#### V. TRANSFER ONTO ROBOTS

To devise robotic grasping strategies that leverage insights captured in the model derived from human trials, we must transfer both the action primitives and the ability to decide on the most appropriate transitions between them [13]. We previously demonstrated that the action primitives of our taxonomy can be implemented on robots [1]. Here we argue that the description of human grasping as sequences of action primitives provides a useful tool to extract transition-based grasping plans for robots. The analysis of human grasping behavior in terms of transition probabilities between primitives allows us to identify the conditions that determine the decision for one path of action over another, and to specify the conditions these transitions depend on.

The results of our data analysis show that human subjects reliably use different action sequences for different object shapes (Section IV-D) and different environmental conditions (Section IV-E). The analysis of failure cases reveals factors that influence the appropriate choice of successor actions (Section IV-F). We used these observations to extract the following list of conditions about the object and the environment that affect the transitioning between primitives in successful grasping. A robotic grasp planner must be able to evaluate the conditions in order to reproduce the grasping behavior observed in human trials. Hence, a robot should be equipped with sensors that provide sufficient information to evaluate the following elementary conditions:

- $\alpha$  The object is in a suitable position and orientation such that when closing the hand the fingers will cage the object.

- $\beta$  The object is flat.
- $\gamma$  The position, the orientation and the shape of the object allow for performing the action primitive 'flip' (f)
- $\delta$  The object is at the edge of the support surface and allows for the action primitive 'edge grasp' (e)
- $\tau$  The hand can establish a force closure with the object
- $\lambda$  Hot surface condition

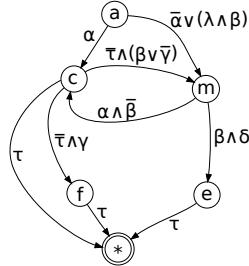


Fig. 7. Decision graph showing statements about the object and the environment that need to be evaluated to enable the respective transition. A transition is only possible when the respective statement is true. If no transition is possible, the primitive fails.

Fig. 7 shows the conditions associated with the transition as logic statements based on the elementary conditions. A transition is activated only when the respective logic statement is evaluated to be true. If no transition is possible, the primitive fails. When multiple statements are true all respective transitions are enabled. In those cases the data obtained from humans was not sufficient to distinguish between alternatives. In these cases, a planner could employ additional heuristics to choose between successor primitives.

We show if all transition conditions can be implemented on a robot a grasp planner program can be designed by following the logic of our decision graph in Fig. 7. Thus, provided that all elementary conditions can be realized with the sensors of a robot, we can transfer—at least in principle—the insights gained from human grasping to robots. Clearly, this transfer depends on a robot hand that is anthropomorphic. To transfer the insights from human grasping to robot hands of different morphology will be part of our future work.

## VI. CONCLUSION

We analyzed 3,400 trials of humans grasping single objects. Our goals in this research is to develop an understanding of the principles of robust human grasping and to transfer them onto robotic systems. Based on our prior work [1], we hypothesized that the exploitation of environmental constraints is a key contributor to human grasping performance. We confirmed and extended our prior evidence in this regard, most importantly by demonstrating that human grasp performance deteriorates significantly when the use of environmental constraints is suppressed. Our analysis also reveals that taking environmental constraints as the structuring principle behind human grasping, we obtain a simple Markov chain description, accurately capturing the grasping behavior observed in all trials. This further supports our

hypothesis that environmental constraints play an important role in human grasping. It also opens up the possibility of transferring human grasping behavior to robotic systems: both the individual grasping primitives observed in humans and the transition conditions between these primitives can be transferred to a robotic system. We therefore advocate environmental constraints as a key concept towards improving the performance of robot graspers.

## REFERENCES

- [1] F. Heinemann, S. Puhlmann, C. Eppner, J. Álvarez-Ruiz, M. Maertens, and O. Brock, “A taxonomy of human grasping behavior suitable for transfer to robotic hands,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4286–4291.
- [2] M. R. 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.
- [3] T. Feix, R. Pawlik, H.-B. Schmiedmayer, J. Romero, and D. Kragic, “A comprehensive grasp taxonomy,” in *Robotics, Science and Systems: Workshop on Understanding the Human Hand for Advancing Robotic Manipulation*, 2009, pp. 2–3.
- [4] I. Bullock, J. Zheng, S. de la Rosa, C. Guertler, and A. Dollar, “Grasp frequency and usage in daily household and machine shop tasks,” *IEEE Transactions on Haptics*, vol. Early Access Online, 2013.
- [5] L. Y. Chang, S. S. Srinivasa, and N. S. Pollard, “Planning pre-grasp manipulation for transport tasks,” in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 2697–2704.
- [6] D. Kappler, L. Chang, M. Przybylski, N. Pollard, T. Asfour, and R. Dillmann, “Representation of pre-grasp strategies for object manipulation,” in *2010 10th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2010, pp. 617–624.
- [7] R. Balasubramanian, L. Xu, P. Brook, J. R. Smith, and Y. Matsuoka, “Physical human interactive guidance: A simple method to study human-planned grasping,” *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 899–910, 2012.
- [8] F. Wörgötter, E. E. Aksoy, N. Krüger, J. Piater, A. Ude, and M. Tamisunaite, “A simple ontology of manipulation actions based on hand-object relations,” *IEEE Transactions on Autonomous Mental Development*, vol. 5, no. 2, pp. 117–134, 2013.
- [9] C. Mandery, J. Borras, M. Jochner, and T. Asfour, “Analyzing whole-body pose transitions in multi-contact motions,” in *IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 2015, pp. 1020–1027.
- [10] M. Kazemi, J.-S. Valois, J. A. Bagnell, and N. Pollard, “Human-inspired force compliant grasping primitives,” *Autonomous Robots*, vol. 37, no. 2, pp. 209–225, 2014.
- [11] C. Eppner, R. Deimel, J. Alvarez-Ruiz, M. Maertens, and O. Brock, “Exploitation of environmental constraints in human and robotic grasping,” *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 1021–1038, June 2015.
- [12] K. Bernardin, K. Ogawara, K. Ikeuchi, and R. Dillmann, “A sensor fusion approach for recognizing continuous human grasping sequences using hidden markov models,” *IEEE Transactions on Robotics*, vol. 21, no. 1, pp. 47–57, 2005.
- [13] D. R. Faria, R. Martins, J. Lobo, and J. Dias, “Extracting data from human manipulation of objects towards improving autonomous robotic grasping,” *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 396–410, 2012.
- [14] D. Kulić, W. Takano, and Y. Nakamura, “Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains,” *The International Journal of Robotics Research*, vol. 27, no. 7, pp. 761–784, 2008.
- [15] M. T. Mason, N. Pollard, A. Rodriguez, and R. Kerwin, “Harnessing human manipulation,” *NSF/ARL Workshop on Cloud Robotics: Challenges and Opportunities*, 2013.
- [16] V. I. Levenshtein, “Binary codes capable of correcting deletions, insertions and reversals,” *Soviet physics doklady*, vol. 10, no. 8, pp. 707–710, 1966.