

Analysis of Human Grasping Behavior: Object Characteristics and Grasp Type

Thomas Feix, *Member, IEEE*, Ian M. Bullock, *Student Member, IEEE* and Aaron M. Dollar, *Member, IEEE*

Abstract—This paper is the first of a two-part series analyzing human grasping behavior during a wide range of unstructured tasks. The results help clarify overall characteristics of human hand to inform many domains, such as the design of robotic manipulators, targeting rehabilitation toward important hand functionality, and designing haptic devices for use by the hand. It investigates the properties of objects grasped by two housekeepers and two machinists during the course of almost 10,000 grasp instances and correlates the grasp types used to the properties of the object. We establish an object classification that assigns each object properties from a set of seven classes, including mass, shape and size of the grasp location, grasped dimension, rigidity, and roundness. The results showed that 55 percent of grasped objects had at least one dimension larger than 15 cm, suggesting that more than half of objects cannot physically be grasped using their largest axis. Ninety-two percent of objects had a mass of 500 g or less, implying that a high payload capacity may be unnecessary to accomplish a large subset of human grasping behavior. In terms of grasps, 96 percent of grasp locations were 7 cm or less in width, which can help to define requirements for hand rehabilitation and defines a reasonable grasp aperture size for a robotic hand. Subjects grasped the smallest overall major dimension of the object in 94 percent of the instances. This suggests that grasping the smallest axis of an object could be a reliable default behavior to implement in grasp planners.

Index Terms—Human grasping, manipulation, activities of daily living, prosthetics, robotic hands

1 INTRODUCTION

THE human hand is a marvelous multifunctional tool, and analyzing its diverse functionality helps both to understand human manipulation behaviors as well as motivate design choices for mechanical hands intending to replicate its abilities. A number of studies have investigated classifying grasps into a discrete set of types [1], [2], [3], [4], [5], [6], and others have been aimed at understanding certain aspects of human hand usage [7], [8], [9], [10], [11]. However, there has been little work on investigating the properties of objects humans interact with on a daily basis and correlating those to grasp choice. A better appreciation of the types and properties of objects that humans commonly manipulate is important in many domains. It can inform hand rehabilitation by focusing on the ability to grasp the most common and important object sizes and masses. Prosthetic and robotic hand designers can improve the performance of their hands by creating a hand that can grasp the most common objects the human interacts with.

Grasp and object properties have previously been studied to some extent by researchers in the robotics, haptics, and psychology communities. For example, it has been shown that the number of fingers used for grasping increases with the size and mass of the object [12], [13] until a two-handed grasp is required, indicating that object size and mass are strong factors in determining the grasp type. In a similar

study [14], it was shown that the contact locations of the fingertips are very similar within and between subjects, with the shape of the object having a large influence on the applied grasp. Klatzky et al. [7] indicated that humans use the same or similar grasp types for certain types of objects, and were able to predict a large percentage of grasps based on the object shape data alone. Researchers in the robotics community have investigated grasp planners based on observation of human behavior [1], [15], [16]. These planners use features such as object mass, local graspable features, and major object axes. It is not clear however, which object property ranges are common in human environments.

In the work described in this paper, we investigate the relationship between grasp types and object properties in a much larger set of human grasping data than has been previously analyzed. We utilize a data set of almost 10,000 grasps taken from a small head-mounted camera from four subjects. For each grasp instance, grasp type, object grasped, and the timing of the grasp are recorded. Prior work by the authors utilizing this data set has analyzed the frequency of each grasp type [17], [18], as well as how to best choose a small set of grasps to span a wide range of objects [19] (part of which is presented). Finally, we made the data set public, including the information on the objects [20].

We classify each object with a set of parameters that are relevant for grasping, taking inspiration from the theory of affordances [21] as we are only interested in the objects that afford grasping and manipulation. In addition to being affected by the properties of the object, grasp choice is also influenced by the task—the relationship between grasps and tasks is presented in the accompanying paper [22], which also investigates the object-task-grasp complex.

This paper is structured as follows. Section 2 presents the object classification that we apply to the data. Section 3

- The authors are with the Department of Mechanical Engineering and Materials Science, Yale University, New Haven, CT.
E-mail: {thomas.feix, ian.bullock, aaron.dollar}@yale.edu.

Manuscript received 12 Aug. 2013; revised 7 May 2014; accepted 19 May 2014. Date of publication 28 May 2014; date of current version 12 Sept. 2014. Recommended for acceptance by A. Okamura.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TOH.2014.2326871




Fig. 1. One object can have multiple natural grasp locations. Depending on the task and other parameters, the human chooses one of the grasp locations to manipulate the object. The only parameter which is on the object level is the mass of the object. All other parameters are individual for each grasp location.

describes the data set on which the object taxonomy was applied and Section 4 presents the results. Finally Section 5 discusses the results and Section 6 concludes the paper.

2 OBJECT CLASSIFICATION FOR GRASPING

2.1 Scope

In order to analyze the relationship between objects and grasp choice, we must first provide a framework into which all commonly manipulated objects fit into. We aim to classify only objects that allow human manipulation. For example large objects, such as a car or a lathe will not be considered as a whole part in this analysis, but will be broken into a set of smaller, more graspable parts.

2.2 Definition of an Object and Grasp Location

For the scope of this analysis, we define an object to be a physical entity that acts as a whole when it is manipulated. Therefore, the actual interaction with the object defines what is regarded as an object. For example, if a stack of papers is manipulated as one, the complete stack is regarded as one object. However, when one grasps only one sheet of paper, the single sheet is the manipulated object. By stacking or dividing objects, a new graspable object can be created.

Most graspable objects do not have any physical connection to the environment, thus making it easy to determine the boundary between the object and the environment. However, some cases are less clear. One prominent example is different bodies connected via joints or hinges. Our classification considers joints or hinges that connect the object to an immobile object attached to the ground (e.g., a wall) to define the “end” of the object. Examples would be a door hinged to a frame or a drawer sliding on a track. However, joints that do not serve to connect the object to the ground are considered to be part of the object (such as the door and door handle or two halves of a pair of scissors). If, during the course of one grasp, the end of an object changes, the larger object is assigned. For example, a closed door is immobile (thus the object would only be the handle) until the handle is turned, upon which it can be opened (in which case the object is the door). In that case the assigned object is the door including the handle, for both sub-movements.

An object can often be manipulated in many different ways. For each way it is manipulated, there might be different proportions of the object relevant for the actual grasp. Therefore, we introduce the concept of a *grasp location*, which

we define as the local part of the object specific to the grasp instance. An object can have multiple grasp locations, and humans will choose a grasp location based on the task and other parameters. For example, a mug (Fig. 1) has at least three different grasp locations: the side (cylindrical shape); the top (disk shape); and the handle (thin curved shape). While it might appear odd to break up an object by grasp locations, most objects have a small set of local geometries that are commonly involved in grasps. This is especially true for objects designed for human interaction. Thus, looking only at the overall object geometry would miss the most relevant parts of the geometry for actual grasping behavior.

2.3 Properties of an Object and Grasp Location

While objects can be assigned any number of a wide range of properties, we place our focus on properties that are important for grasping and that can be easily assessed. The right side of Fig. 1 shows an overview of object properties that we consider, which are explained in more detail below. The only property defined on the full object level is the mass of the object. The other parameters are defined for each grasp location, which can have very different local shapes.

The *mass* is a mostly straight-forward parameter to assign. For the purposes of this classification, we choose to define mass irrespective of the force needed to manipulate the object, and we explicitly take the full mass of the object into consideration. For example, even though the large mass of a door is supported by its hinges, we record its full mass for completeness. In our view the task ultimately defines how the mass of the object relates to the grasp force and is therefore taken into account in our task taxonomy [21]. Depending on if the door is transported or simply opened, the resulting grasp force can change to a large degree. For other objects, the mass is variable, such as in a bottle of water. For the purposes of our classification, if the fill level is unknown, as is frequently the case when visually determining object properties, then a typical fill level of around half of the container’s maximum capacity was assigned. For very heavy objects or objects that are completely fixed to the environment, we bin the mass into a single category (1,000 + g) for the purposes of the analysis presented in this paper.

Besides mass, all the other object properties that we assign are specific to the individual grasp locations. The next three properties define the *shape and size* of the grasp location. After much deliberation as to how to distill such




Fig. 2. Grasped dimensions for two objects. In the case of the round object, either of the two diameters could define the grasped dimension. Thus the two dimensions a and b of the object are both assigned to be the grasped dimension.

a massive variation in shape down to a manageable number of parameters, we decided to employ an established nomenclature that classifies convex objects according to their three primary axes [23], [24]. For objects of ellipsoid or convex shape there are four distinct shape classes, which are summarized in Fig. 3, left column. The three major axes of the object define the major directions of the object. These axes are denoted a, b, and c, where a is along the longest object dimension and c is along the shortest. We refer to length of the object in those three dimensions as the size of the grasp location (A, B, and C), where $A \geq B \geq C$. The constant R defines the threshold above which two axes are considered to be different lengths (i.e. if the ratio between two sides is greater than R , they are considered to have a different length). Tests determined that the value at which one typically regards two axes to be different is about $R = 3/2$ [23]. Although these shape categories were conceived for ellipsoid shapes, the general concept can be applied for all objects. It has been argued [23] that the assignment of those three values can be done with an

agreement between raters of less than 5 percent, and our observation confirms that. There are instances where the long dimensions are much longer than any human hand. In this case it does not add information relative to grasping to precisely record the long dimension, since they are too large to be grasped. Therefore, any dimension larger than the largest thumb-index finger span (approximately 15 cm [25]) is binned into the 15+ cm category.

We define the *grasped dimension* as the part of the object that lies between the fingers when grasped. By referring to the previously defined object axes, the values are from the prior set {a, b, c} to indicate which axes best determine the hand opening. Fig. 2 shows two examples of how the dimensions determine the grasped dimension. In the left figure, the object is grasped along the shortest dimension, therefore the grasped dimension is c. For many objects, more than one dimension determines the hand opening, for example for a disk (as shown in Fig. 2, right image) the length A and B are equal and both determine the hand opening. Therefore, the grasped dimension for Fig. 2, right, should be set to a/b. That means that both the a and b axis determine the hand opening. In particular for such a round object, there is no unique labeling, as the a and b axes can be rotated around c without change in length.

Many common objects have circular or elliptical cross sections. Therefore, we introduce the parameter *roundness*, which defines whether the object is curved around certain axes. Roundness can be expressed along the three previously defined axes. The possible values of roundness are any combination of {a, b, c}. Roundness is assigned to a given {a, b, c} axis if the object cross sections perpendicular to that axis are circular or elliptical. In the case of Fig. 2, right, there is roundness around the c axis. In the case of a sphere, a cross section in any direction is circular, thus the roundness parameter is “abc”. Fig. 3 gives an overview of the possible roundness values and gives example images.

| Zingg's object categories | Roundness along axis | Object Type |
|---|----------------------|----------------------|
| Equant $B \leq A < R \cdot B$ $C \leq B < R \cdot C$ | - | cubic (1) |
| | a b c | cubic cylinder (2) |
| | abc | sphere (3) |
| Prolate $A > R \cdot B$ $C \leq B < R \cdot C$ | - | long prism (4) |
| | a | cylinder (5) |
| | b c | long ellipse (6) |
| | abc | long ellipsoid (7) |
| Oblate $B \leq A < R \cdot B$ $B > R \cdot C$ | - | short prism (8) |
| | a b | short ellipse (9) |
| | c | disk (10) |
| | abc | short ellipsoid (11) |
| Bladed $A > R \cdot B$ $B > R \cdot C$ | - | irregular (12) |
| | a | ellipse A (13) |
| | b | ellipse B (14) |
| | c | ellipse C (15) |
| | abc | ellipsoid (16) |




Fig. 3. Object types that can be inferred based on the object dimensions and the roundness. The inequalities define into which object category of Zingg [23] the object belongs to. We followed his recommendation and set $R = 3/2$.

By looking at the relative length of the grasp location dimensions, one can infer the shape categories defined by Zingg [23], [24]: equant, prolate, oblate, and bladed (Fig. 3 left). By adding the information about the roundness of the object, that categorization can be further refined. For each category there are five possible roundness conditions: no roundness, roundness along each individual axis or roundness along all axes. This leads to a set of 20 combinations, with an overview in Fig. 3. Since in some categories the object dimensions are equal, they cannot be distinguished and thus different roundness values around them are meaningless. If, for example, all dimensions are equal (equant category), roundness around A is equal to roundness around C, simply because the labeling of the dimensions is not unique. By merging those ambiguous categories, the possible combinations are reduced down to 16 distinct object types. Some of the object types resemble common objects, such as cylinders or disks while others have no known common name.

The final parameter, *rigidity*, defines how the object responds when force is applied by the grasp. There is evidence that the human adjusts the grasp parameters accordingly [26]. We assign four possible properties: {rigid, fragile, squeezable, floppy}. A grasp location is defined as *rigid* if it can withstand the full hand force without breaking or deforming significantly. For example, a mug or a piece of wood would be regarded as rigid. The *fragile* category also does not show major deformation when force is applied. However, there is a threshold, upon which the object will break or start to deform a lot. Examples for this category are an egg or a very slender glass. The next category is *squeezable* objects. These deform significantly under normal grasping forces. Examples are sponges or paper cups. Even though these objects are easy to deform, they still have an inherent shape, which distinguishes them from the last category. Finally, an object is *floppy* if it has no inherent shape, such as a rag. Those objects deform heavily under gravity. Because the object shape can change so much, size, grasped dimension and roundness are not assigned for floppy objects.

2.4 LIMITATIONS OF THE CLASSIFICATION

To fully describe real objects many parameters are needed. The presented classification incorporates important object properties that are likely to govern the choice of grasp. Many other properties, however, were not incorporated into the classification, some of which are discussed in the following.

The grasp force is influenced by many factors—currently only the effects of object mass are included. Among other things, the center of gravity and its distance from the grasp location is not taken into account. If the grasp point is far away from the center of gravity, a high grasp force could be required even for a light object. The surface friction is also not recorded. We did not include these parameters as those parameters are too difficult to obtain in an unstructured environment with current technology. Furthermore, many of these parameters are related to the task being done with the object, and we add detail along those lines in the accompanying paper [22].

There are other additional parameters that would afford or disallow manipulation. Very hot or cold objects would require different manipulation strategies as they would

otherwise injure the hand. A knife disallows grasps that would touch the cutting edge. Because these parameters are only applicable to a small subset of objects or are hard to measure, we exclude them from our analysis. Furthermore, we only analyze instances where the object is grasped, thus we assume implicitly that every object-grasp combination has to be affordable; otherwise it would not be observed by our methodology.

3 METHODS

3.1 Experimental Procedure and Apparatus

Two machinists and two housekeepers participated in the study discussed in this paper. Those two professions were selected because both perform usually manual labor. Furthermore, the housekeeper works at home, which can be relevant for robots that operate in home environments. The machinist represents manufacturing, where there are many different tools and machines present. Both areas are important areas in which robots are supposed to operate. We used two subjects per profession to have a minimum amount of statistics.

The following enrollment criteria were used to screen potential subjects for the study: significant experience as professionals in their field, of normal physical ability, right-handed, able to participate for long enough to generate eight hours of data, and performing tasks representative of their profession during the span of their participation. The protocol was approved by the local IRB.

Full details of the experimental protocol can be found in [18], but a brief summary follows. Subjects wore a head-mounted camera that recorded their hand use during normal professional work, for at least eight hours per subject. One rater then tagged the right-handed grasps in the video, including the high level task name the subject was performing and the object that was grasped. Examples for a task name would be “wiping” and for the object “spray bottle”, “mop” and “sponge”. The reliability of the tagging process was evaluated by sections of the videos that were tagged by two raters. Agreement of the raters was found to be Cohen’s $\kappa = 0.54$ [18].

3.2 Applying the Object Classification

The final object classification was derived after the video data was recorded and tagged with regard to grasp type, high level object and task name. Reviewing the wide range of grasping behavior through the video data and generated object snapshots informed the development of the classification.

The object categorization builds on the initial classification as presented in our previous publication [18]. We make use of the grasp classification and add the object classification layer. The most common high level object names can be seen Fig. 6. This object name data was then used to classify objects accordingly. Because the name does not allow distinguishing grasp locations, we reviewed generated grasp snapshots for each object to pick a representative, most common grasp location for each object. Therefore, we will use grasp location and object interchangeably in the remainder in the paper. The precision of the assignment could be improved if each grasp instance would be reviewed independently and the corresponding grasp location assigned, however that would need

TABLE 1
The Table Shows the Inter-Rater Agreement
for the Object Classification

| | Agreement [%] | Cohen's kappa | Pearson Correlation |
|-------------------|------------------|------------------|------------------------|
| Cannot Classify | 86 | 0.47 | |
| Grasped Dimension | 66 | 0.52 | |
| Rigidity | 81 | 0.63 | |
| Roundness | 81 | 0.67 | |
| Size A | | | 0.78 |
| Size B | | | 0.86 |
| Size C | | | 0.91 |
| Grasp Size | | | 0.77 |
| Mass | | | 0.64 |

Both the overall agreement and the Cohen's kappa [36], respective the Pearson correlation [27] are reported.

a complete retagging of the data set, a significant amount of work. We therefore think this intermediate step via the high level object name is a good compromise.

Overall, different instances of the same high level object were relatively consistent. This is mainly because a given object was handled similarly multiple times and many objects only had a few instances. Naturally, high level objects with few instances are easier to classify accurately since the diversity of the objects within the group goes down. To accommodate object classes that are very diverse, the raters are given the choice of not classifying an object. By not considering these object classes that are very broad, the overall precision of the classification can be improved. For example, the "lid" and "small box" object descriptors were too broad. Although "pan" was relatively specific, it was grasped partly by the handle and partly by the pan itself, so this variation in grasp location precluded a meaningful classification.

A list of measured and weighted benchmark objects was presented to the raters in order to better allow them to estimate size and weight from the video. It was not possible to measure the real objects, as the object classification was derived sometime after the video recording and we did not have access to all four locations. This step would increase tagging accuracy, at the cost of a largely increased effort. Furthermore, the accuracy of the classified data could be increased if every grasp instance in the data set was assessed individually. However, we believe that our approach of picking representative instances is a good compromise between tagging effort, which already took hundreds of hours of work, and precision of the results. The current precision can be tested with inter-rater statistics.

3.3 Inter-Rater Agreement

On the full set of 406 objects, the two raters agreed 86 percent of the time (Cohen's $\kappa = 0.47$) on whether each object type can be classified. Most of the "cannot classify" objects were due to large variability within the object class. For the further inter-rater analysis, all objects rated as "cannot classify" by at least one rater were discarded. For this analysis this is necessary as data from both raters has to be present in order to compare them.

The nominal classes can be directly compared using Cohen's κ , and their inter-rater agreement values are

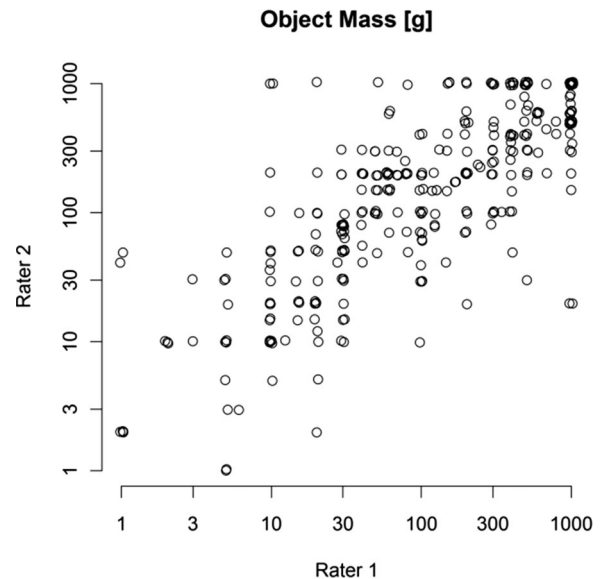


Fig. 4. The figure shows the mass classification of the two raters. Objects that are fixed with the environment or that have more than 1,000 g mass were reduced to 1,000 g. Jitter was added to aid the visualization and the plot uses a logarithmic scale.

reported in Table 1 For the numerical values a Pearson correlation [27] is shown.

Concerning the grasped dimension, the biggest disagreements were rater 2 assigned b/c and rater 1 assigned c (30 objects) and where rater 2 assigned b/c and rater 1 b (19 instances). Those 49 objects make up a large part of the disagreement, however as those are "neighboring" parameters the resulting error in the classification is less than the low classification rate would suggest. Overall, rater 2 had a stronger tendency to assign b/c, he assigned this property for 55 more objects than rater 1.

The confusion matrix for rigidity is relatively symmetric, indicating that there was no systematic difference between both raters.

For the roundness parameter the only common confusion is between no roundness and roundness along a. The confusion is symmetric, indicating no systematic bias.

The rater agreement for the object mass is shown in Fig. 4. The plot shows that there is considerable disagreement between raters, which is also reflected in the lower Pearson correlation coefficient of 0.64. However, the general trend holds that a heavier object is assigned a larger mass. Objects heavier than 1,000 g or those that are fixed to the environment were set to 1,000 g. That explains the lines present at 1,000 g for both raters. In particular, rater 2 had the tendency to set objects, like handles and knobs of machines, which were connected via a hinge to the environment to be fixed to the environment. However, the classification requires setting those objects to their actual weight, as they are commonly handled as an individual object. This explains many of the instances where rater two is at 1,000 g and rater one has assigned a lower mass value. On the other side, the bound at 1,000 g of rater 1 is more due to different opinions about the object weight.

The agreement on the grasped dimension, as shown in Fig. 5 shows a consistent trend in the assignment. Note that the grasp size is a combination of the assigned object sizes and the grasped dimension. Therefore this plots tests the




Fig. 5. The figure shows the grasp size as determined by the two raters. The grasp size is the length property of the object that best defines the hand opening. A jitter of 1 mm was added to the data to help visualization. Note that the grasp size is a combination of the size assignment and the assignment of the grasped dimension.

precision of the size assignment as well as the assignment of the grasped dimension.

The correlation of the mass is lower than for the size values. The reason for this is likely due to the applied methodology—the mass is estimated based on the video recording, and the size of the object is more directly observable. The sizes B and C are the most relevant dimensions for grasping and have the highest correlation values.

3.4 Estimation of the Error

All the 406 objects in the data set were classified by two different raters. The object properties outlined in Section 2 served as a guideline for the raters on how to classify the objects.

The error bars in the figures are calculated using a case resampling bootstrap method [28]. Random samples are taken from the original data set, with replacement. The number of samples matches the length of the real data set. Furthermore, the object properties are taken randomly from one of the two raters' classification data sets. Both random samplings are based on a uniform distribution. The data was resampled 1,000 times and 95 percent confidence intervals were calculated based on the standard deviation of the results. The reported mean corresponds to the data set that is based on the merged classification of both raters.

3.5 Final Object Classification Data Set

The classifications of the two raters were combined in a semi-supervised method. The differences in the parameters were analyzed by rater 1, who made a final decision about which rater's assignment was correct. This step allowed correcting obvious classification errors, and some of the differences revealed valuable information about limitations of the classification. For the numerical parameters (size and mass), the mean was taken. However, in particular for the mass, some values were corrected—mostly objects that were

incorrectly assigned to be fixed to the environment were set a real mass. The error bar estimation reflects the disagreement of the raters, thus gives a clear estimate on how reliable the results are.

Note that the data set is not identical to [18] as we discard instances that could not be classified and we only analyze frequencies, whereas [18] also looks at the duration of grasps. The full data set is available online for download [20].

3.6 Data Processing

Following the definition of the object properties, the maximum object size bin was set to "15+ cm", and all larger values were included in it. Accordingly, in the a, b, and c dimensions, 3,830, 497, and three (42, 5 and 0 percent of the data set) instances, respectively, were set to 15+ cm.

The mass of all fixed objects, as well as all objects heavier than 1 kg, were set to 1,000+ g. Only 475 instances are heavier than 1,000 g and two instances are fixed to the environment. Those two categories combined make up 5 percent of the data set. By setting this weight limit, we set the focus of our analysis onto the lower weights, which are more common. Furthermore, 1,000 g corresponds to a typical transition point from a one handed grasp to a two handed grasp [13]. As the taxonomy used for defining the human grasp types is focused on one handed grasps, we think this transition point offers a natural point for the cutoff.

4 RESULTS

4.1 General Statistics

After preprocessing the data, the remaining 9,100 grasp instances and 306 objects are used for further analysis. This number is lower to [18], as some instances are removed because raters could not classify the object properly. The first few objects are very common, about 2/3 of the objects have less than 10 instances.

Fig. 6 shows the first 12 objects and their distribution across the grasp types, and Table 2 shows their properties. The most common grasp type is the medium wrap with 1,339 instances in the data set. The "mop" and "spray bottle" objects are prevalent in that grasp type, accounting for almost 50 percent of the data for that grasp. Precision disk is heavily weighted by two objects, "towel" and to a lesser degree, "sponge". Therefore, those two objects will dominate the object properties handled by this grasp type. Also, many other grasp types are heavily influenced by only a few objects. These common objects will potentially distort the apparent characteristics of the grasp, rather than giving a more comprehensive assessment of the grasp. However, most grasps are used with many different objects and should produce reliable results.

4.2 Overall Object Properties

Fig. 7 shows the statistics for the nominal values of the data set. Object rigidity is dominated by rigid objects, which make up about 58 percent. The *fragile* category is only 0.2 percent of the data set. Squeezable objects are 18 percent of the data set and floppy objects are 24 percent.

For roundness, the data shows that a large proportion (61 percent) of the non-floppy objects show roundness along the long axis. Non-round objects make up 36 percent, whereas

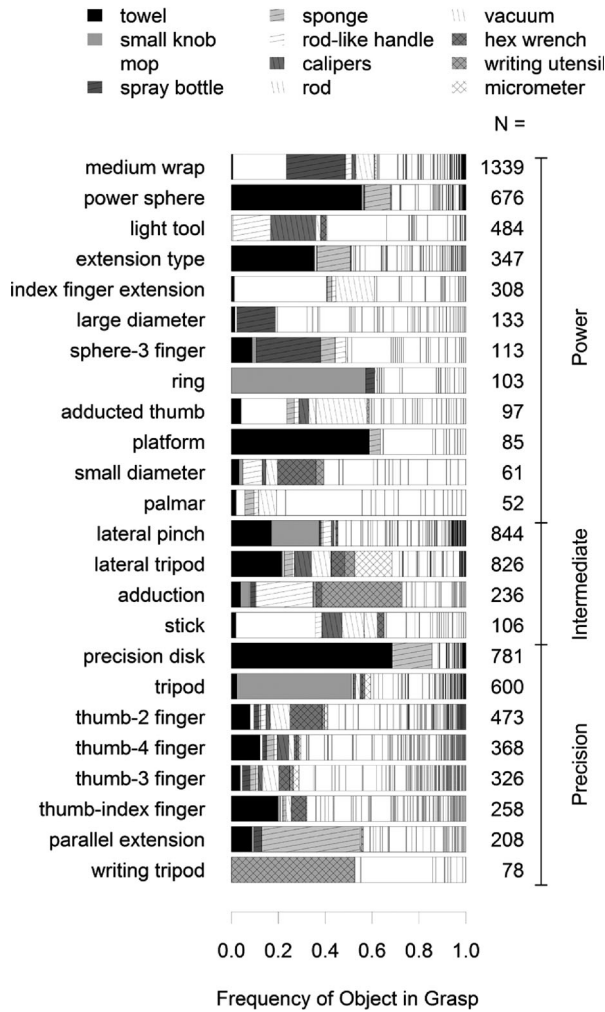


Fig. 6. General object statistics. The most common 12 objects are highlighted, all other objects are only indicated in white separated by vertical lines. One can see that some grasps are heavily weighted by some of the most common objects. Grasps with fewer than 50 instances are not plotted. Total number of instances is 9,100. The order of the legend corresponds to the order within each bar.

other symmetries make up for only 3 percent in the non-floppy objects. Many objects in the data set are cylindrical and grasped from the side, having “b/c” as the grasped dimension.

In about 94 percent of the grasp instances the smallest dimension is equal to the grasped dimension (a/b/c, b/c, or c). In 1 percent of the cases the grasped dimension is a/b and in 4 percent of the cases it is b.

Concerning object size (Fig. 8), it is interesting to see that for A, the largest dimension, values of 15+ cm dominate the data set. This is the case for 55 percent of the non-floppy objects. In all these cases the longest dimension cannot be grasped, since it is larger than the human hand opening. The second largest dimension, B, shows much smaller values. For B, only 7 percent of the sizes are too large to be grasped and 89 percent of the dimensions are even smaller than 7 cm. This means that most sizes are already well within the graspable hand span (note that the non-prehensile “platform” grasp does not require a small object dimension to execute). Concerning C, in 94 percent of the cases the dimension is smaller than 5 cm.

As shown in Fig. 9, the mass of the objects is relatively evenly distributed up to 400 grams. Heavier objects are

TABLE 2
The Table Shows the Most Common 12 Objects and Their Properties

| Count | Name | A [cm] | B [cm] | C [cm] | Grasped Dimension | Rigidity | Roundness | Mass[g] |
|-------|-----------------|--------|--------|--------|-------------------|----------|-----------|---------|
| 1645 | towel | - | - | - | - | f | - | 250 |
| 550 | small knob | 5 | 3 | 3 | b/c | r | a | 150 |
| 522 | mop | 15+ | 4 | 4 | b/c | r | a | 400 |
| 451 | spray bottle | 15+ | 6 | 4 | b/c | s | - | 400 |
| 440 | sponge | 10 | 6 | 3 | c | s | - | 17.5 |
| 259 | rod-like handle | 15+ | 1 | 1 | b/c | r | a | 400 |
| 232 | calipers | 15+ | 3 | 1 | b/c | r | - | 171 |
| 206 | rod | 15+ | 1 | 1 | b/c | r | a | 300 |
| 200 | vacuum | 15+ | 5 | 5 | b/c | r | a | 1000+ |
| 190 | hex wrench | 10 | 1 | 1 | b/c | r | a | 30 |
| 187 | writing utensil | 15+ | 1 | 1 | b/c | r | a | 20 |
| 174 | micrometer | 12 | 2 | 2 | b/c | r | a | 170 |

The Count column indicates the number of instances that object was found in the data set. All of those objects could be classified accordingly. The abbreviations in the Rigidity column are: f . . floppy, r . . rigid, s . . squeezable. Size values were rounded to the next integer.

rare; only about 5 percent of the objects are heavier than 1,000 g. Furthermore, objects completely fixed to the environment almost never occur, with only two instances in the data set. The two large error bars at 100 and 400 g can be explained by the towel object, which is very common, with 1,645 instances. Rater one assigned the towel 100 g, while rater 2 assigned 400 g. Therefore, depending on which rater value is used, a large number of instances are shifted between those two bins.

Table 3 gives an overview of the distribution of the object types in the data set. The most common objects are cylinders (4,151 instances), irregular objects (1,807 instances) and short prisms (573 instances). From all the possible objects types, almost all have at least some instances. If the number of instances is 0, but there is a range present, this means that one of the raters had such an object assigned, but it was not included in the merged data set.

4.3 Objects by Profession

The two professions work in very different environments. The housekeeper manipulates household and cleaning objects in a home environment. On the contrary, the machinist works in a highly specialized environment with fewer everyday objects. As already shown in [18], the machinist uses a more diverse set of grasp types, and overall it seems that the tasks require a higher degree of dexterity. A major difference from the housekeeper data set is the presence of controls, lever arms and other objects that are connected via a hinge to the environment. These are commonly used to control lathes, drills and other machine tools.

The overall statistics for the machinist and housekeepers are shown in Figs. 7b and 7c respectively.

In the machinist data set, 59 percent of the instances are rigid cylinders, of which all are grasped with a grasped dimension of b/c. The objects that are cylinders range from handles of tools, controls of machines, to drill bits. Rigid

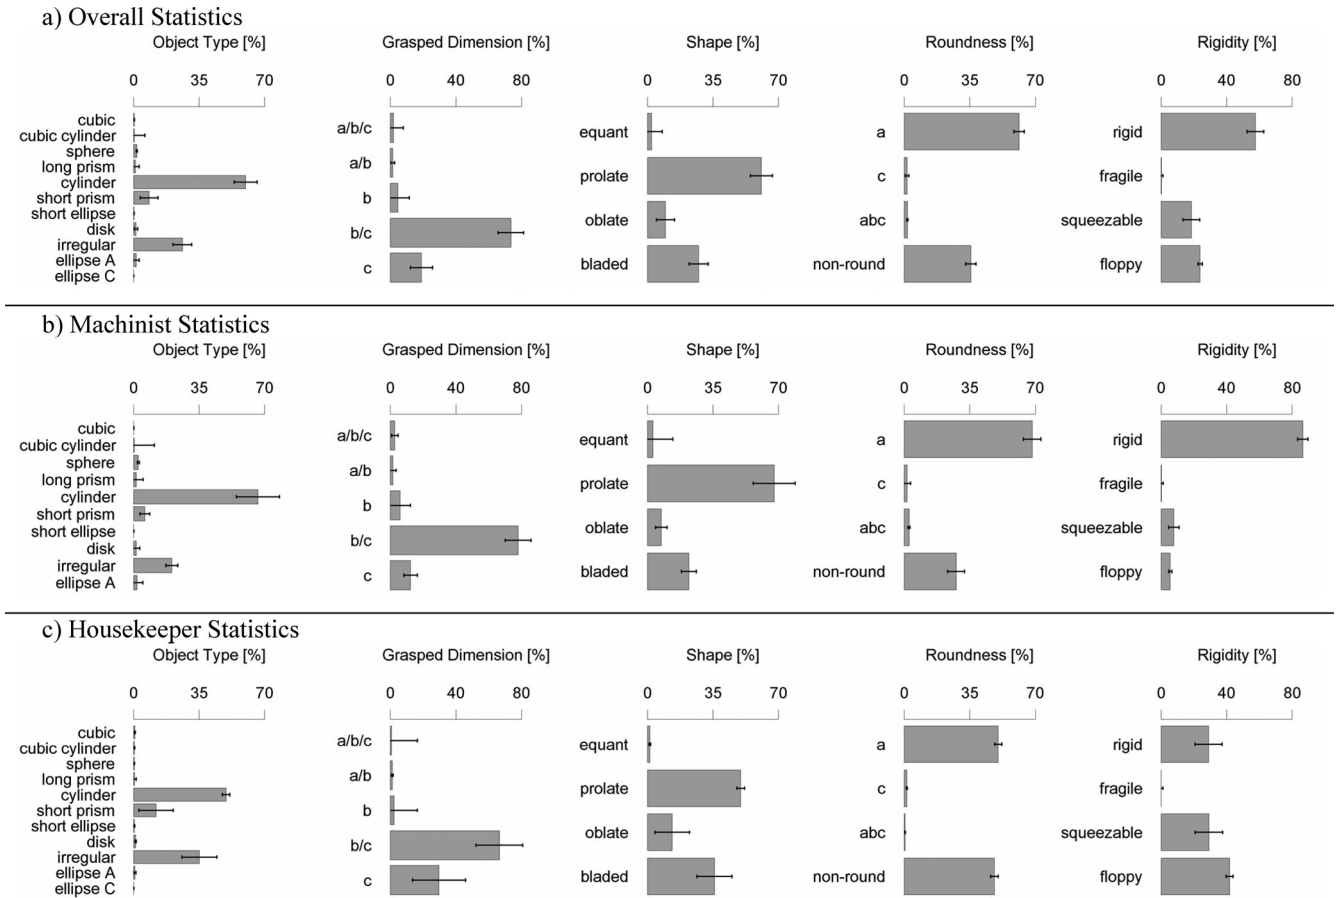


Fig. 7. The figure shows the overall object statistics as well as the statistics per profession. The errorbars correspond to a 95 percent confidence interval; the exact calculation is described in Section 3.4. Note that floppy objects are only shown in the Rigidity category. The other properties are not assigned for floppy objects.




Fig. 8. The figure shows the histogram of the object size. The errorbars represent the 95 percent confidence. Object sizes larger than 15 cm were reduced to 15 cm. Note that by definition $A \geq B \geq C$ and that floppy objects were excluded which reduces the instances to 6,938.

objects are present in 86 percent of the instances and floppy and squeezable objects are present in the machinist data with 6 and 8 percent, respectively. In terms of the grasped dimension, it is dominated by b/c (78 percent) and the other properties are prevalent in 1–12 percent of the non-floppy instances.

The housekeeper data set is heavily weighted by floppy objects, which make up 42 percent of the instances. Furthermore, rigid and squeezable objects are also prevalent in about 29 percent of the instances in the housekeeper data. The large error bar of the rigid and squeezable categories can be explained by the fact that rater 1 and 2 disagreed whether spray bottle (9 percent of the housekeeper data) is rigid or squeezable. Depending on that assignment a large proportion of the data is shifted between rigid and squeezable. Cylinders are slightly less common as in the machinist data set. In terms of the grasped dimension, both b/c and c

are common, in 66 and 30 percent of the non-floppy objects. The other grasped dimensions appear in at most 2 percent of the instances.

For the machinist data set, there are 30 different parameter combinations of grasped Dimension, shape, roundness




Fig. 9. Histogram of the object mass data. Object with mass greater than 1,000 g and that are fixed to the environment are set to 1,000 g. Full data set with 9,100 instances is shown. The errorbars show the 95 percent confidence interval.

TABLE 3
Number of the Instances of the Shape Classes in the Data Set

| Zingg's | Number | Object Type | Number Instances |
|---------|--------|-----------------|------------------|
| Equant | 1 | Cubic | 17 ± 13 |
| | 2 | Cubic Cylinder | 21 ± 511 |
| | 3 | Sphere | 116 ± 27 |
| Prolate | 4 | Long Prism | 68 ± 179 |
| | 5 | Cylinder | 4151 ± 559 |
| | 6 | Long Ellipse | 0 ± 3 |
| | 7 | Long Ellipsoid | - |
| Oblate | 8 | Short Prism | 573 ± 428 |
| | 9 | Short Ellipse | 7 ± 17 |
| | 10 | Disk | 86 ± 87 |
| | 11 | Short Ellipsoid | - |
| Bladed | 12 | Irregular | 1807 ± 454 |
| | 13 | Ellipse A | 91 ± 149 |
| | 14 | Ellipse B | - |
| | 15 | Ellipse C | 1 ± 3 |
| | 16 | Ellipsoid | 0 ± 3 |

The number of instances and the width of the 95 percent interval is shown. The shapes are shown in detail in Fig. 3. The table only contains non-floppy objects (6,938 instances).

and rigidity, whereas for the housekeeper this number is 41. However, to cover 90 percent of the instances, the machinist data set needs the most common eight combinations and the housekeeper needs only six.

4.4 Object-Grasp Type Relationship

The data set can be used to analyze the relationship between the grasp type and object properties.

Fig. 10 shows the distribution of the grasp size (the object dimension which defines the hand opening) and the mass distribution for different grasp types. Violin plots [29] are used to show detailed density estimates for each grasp, due to the irregular shapes of the underlying distributions. We use the R [30] implementation 'vioplot' with standard parameters. The dot represents the median of the data, and the thick line shows the 25th and 75th percentile. The number on the right shows the number of samples for each grasp. On top, the properties of power, intermediate and precision grasps [4], [5], [31] are shown. An intermediate grasp requires both power and precision. The label next to the grasp name indicates which of these three classes the grasp belongs to.

The first three rows of the plots in Fig. 10 show the distribution of the grasp size and mass for the three high level grasp classes. Power grasps are generally applied to large and heavy objects. In terms of the mass and grasp size, intermediate and precision grasps seem to be used for similar objects.

4.5 Grasp Selection by Object Type

Usually humans have a clear idea how an object type is grasped. For example, one generally assumes that a cylinder would be grasped from the side. In order to quantify this

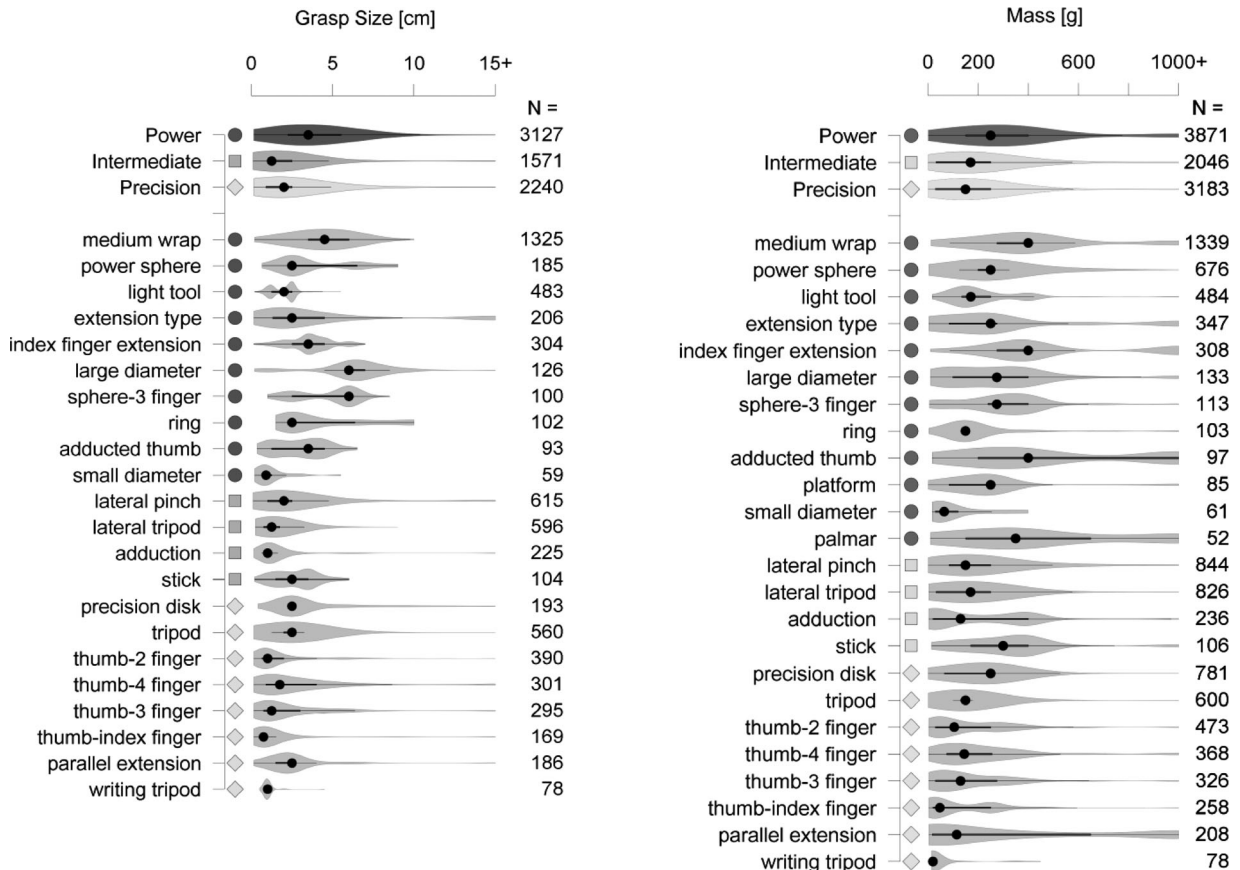


Fig. 10. Grasp Size and mass distribution for each grasp type. The average distribution for power, precision and intermediate grasps is shown on top. The symbols next to the grasp name indicate to which group on top each of the grasps belong to. For example the round symbols indicate a power grasp. Grasp types with fewer than 50 instances are not plotted. Total number of instances for grasp size is 6,938 (floppy objects are excluded) and 9,100 for mass.

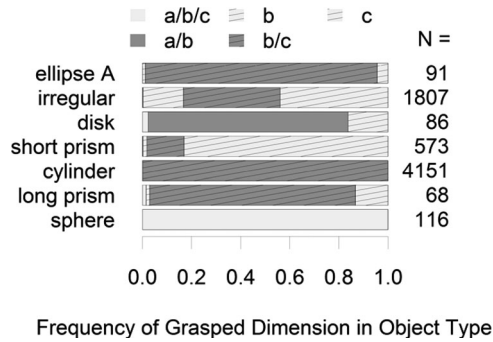


Fig. 11. Comparison between the object type and the grasped dimension. Uncommon combinations with less than 50 instances are not plotted. Floppy objects are excluded which reduces the number of instances to 6,938. The order in the legend corresponds to the order within each bar.

relationship, Fig. 11 shows how object type relates to grasped dimension. Most object types have a clear preferred way to approach them. Long objects are usually grasped from the side, rather than using the longest dimension. A disk is usually grasped from the top (grasped dimension a/b), whereas a short prism is grasped by the smallest dimension. Irregular objects unsurprisingly show the largest variation in the grasped dimension.

4.6 Grasp Use over All Objects

One interesting way to utilize the results presented in this paper is to examine the smaller subset of grasp types that can be used to grasp the most objects, given the fact that many grasps can be used for effectively the same purpose. For example the “mop” object (see Fig. 6) is grasped with a medium wrap, index finger extension, adducted thumb, and stick grasps. By looking over all combinations of grasps and objects, we can assess the versatility of a subset of grasps at handling the largest number of objects in the data set, and calculate a “grasp span” score for each. A grasp set with a high span score is a versatile set which is able to handle a wide variety of different objects. This methodology is explained in more detail in [19].

Fig. 12 shows these results, with the optimal set of grasp types for a given number of allowable grasps in the subset and the corresponding fraction of the number of objects in the set that the set is used for (“Grasp span score”). The results show that the most important grasp type is the medium wrap, which is also the most frequent grasp in the data set, followed by the lateral pinch and the thumb two-finger grasp. Adding successive grasps yield a decreasing benefit in the grasp score where the score shows an asymptotic behavior and reaches 1 for the full set of all grasps. This asymptotic behavior is expected from the grasp span, since the first few grasps ought to be most important in determining the overall functionality of a hand. Because of this effect, results only up to five grasps are presented. While the choice of the first two or three grasps is quite important, after that it is possible to exchange lesser important grasps and achieve a very similar span score.

5 DISCUSSION

5.1 Overall Object Properties

Overall, the object property data can be useful in a number of areas, such as robotic hand design and testing, defining

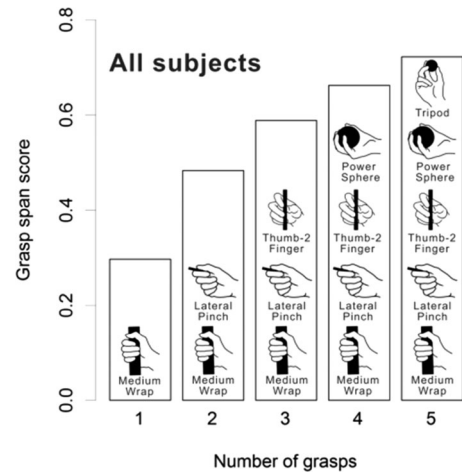


Fig. 12. Optimal grasp sets for any given number of grasps. Adding more grasps gives diminishing returns and already five grasps are close to the maximum score of 1. Medium wrap and lateral pinch are shown to be particularly important grasps overall.

rehabilitation goals and tasks, and designing haptic interfaces or other devices the hand interacts with.

For the objects within this data set, the object properties vary surprisingly little. The objects are relatively lightweight, with about 92 percent of the objects below 500 g and in only about 5 percent of instances the grasped object heavier than 1,000 g. This fits related work showing that the majority of objects in lifting activities were less than 300 g [32] and to [33] who compiled a list of everyday objects that pose a challenge for patients with motor impairments. In a three finger tripod grasp, a female hand can produce an average of 15.7 pounds (7.1 kg) of force with their non-dominant hand [34]. Even for a very slippery object, the hand should be able to lift most of the objects with ease. However, the task based force requirement is not taken into account. One can only speculate on how the task will change the force requirement, but at least for pure transportation and object lifting, the human hand is rarely operated near its maximum strength. In the complementary paper [21] it is shown that in about 60 percent of instances the grasp force is directly correlating to lifting the object. Therefore, in 40 percent of the instances the grasp force is dictated by the task, rather than the mass of the object. This implicates force requirements for artificial hands as well as peak forces needed for feedback in haptic devices. Also in hand rehabilitation it might be beneficial to focus on objects lighter than 500 g.

A similar observation can be made for the grasp size of the objects. The hand is rarely challenged to produce a large hand opening to accommodate the object, likely due in large part to the objects being designed to accommodate grasping. In 5,747 of the cases (83 percent of the non-floppy objects), the grasp size is less than 5 cm, and in 6,827 instances (98 percent of the non-floppy objects), it is opened less than 10 cm. This suggests a natural hand aperture range for robotic hand designs, as well as a range for test objects to determine the grasping capabilities of a particular design. This also suggests that devices designed for the hand, such as haptic devices or other tools, may be most comfortable if they require a hand opening of less than 5cm.

The grasps that handle the objects with the largest size (Fig. 10) are, “large diameter”, “sphere-3 finger” and

“medium wrap”. Interestingly, those grasps are not strongly biased toward heavy objects; only the medium wrap grasp is in the highest mass range (Fig. 10). On the other side of the spectrum, the smallest objects are handled with “thumb-2-finger”, “thumb-index finger”, “adduction” and “writing tripod”. All these grasps have a very low median grasp size of about 1 cm. It seems that for each grasp type there is a relatively narrow mass and size range they are used for.

5.2 Grasped Dimension

The grasped dimension is related to both the task and the object. We believe it fits into the object classification better, as for each grasp location the actual grasped dimension seems to be fairly stable. Furthermore, the grasped dimension is defined in conjunction with the object dimensions, which are inherently linked to the object properties.

As shown in Figs. 7 and 11, subjects have a clear tendency to grasp objects by the smallest dimension of the object. In about 94 percent of objects grasped (excluding floppy objects), the smallest dimension was the relevant grasped dimension, being either a/b/c, b/c, or c. This confirms the results found in another study where elliptical shaped cylinders were grasped. In that case the subjects grasped the minor axis in 67 percent of the trials [35]. This bias seems to be at least partly based on the restrictions the object size imposes. For the largest dimension a, 84 percent of the instances are larger than 10 cm, which already needs a relatively large hand opening, thus might not be favored by humans. Both b and c dimensions are usually much smaller. Only 8 and 0 percent of the objects are larger than 10 cm for the b and c dimension respectively. This data shows our 15 cm cut-off hides little detail of grasping behavior. Thus, it is typically feasible to grasp either of these dimensions, from the perspective of hand opening size. These results suggest that robotic hands and grasp planners should be designed to facilitate gripping the smaller dimensions of common objects, and generally avoid grasping by the largest object dimension. This heuristic seems to make intuitive sense with many common objects.

5.3 Object Types

As summarized in Fig. 3 and Table 3, those types that have established common names (disk, cylinder etc.) are more frequent than the object types that do not. That shows that the object type assignment is aligned with the general perception about objects categories.

The object types also seem to have very specific grasping patterns associated with them (Fig. 11). Long objects are usually grasped in a wrap grasp where the grasped dimension is b/c. A disk, as one would expect, is usually grasped along the a/b dimension, or from the side in a few rare instances. The strong relationship between object types, grasped dimension and the expected intuitive results show, for example, that some basic heuristics may be sufficient to use “human-like” grasping strategies in robotic systems. These results could also be used to better design devices where the intended grasp orientation is clear to the user.

5.4 Most important Grasp Types

The fact that the most frequent type of object is a cylinder fits well to the result of the *grasp span* analysis, where the most important grasp type to handle most of the objects is the medium wrap (Fig. 12). The median mass it handles is 400 g with a grasp size of 4.5 cm. The second most important grasp, the lateral pinch, is used on smaller and more lightweight objects (median mass 150 g, median grasp size 2 cm). The third most important grasp, the thumb-2 finger grasp, adds functionality by grasping even smaller and more lightweight objects (median mass 105 g, median grasp size 1 cm). This small set of versatile grasps, especially the medium wrap and lateral pinch grasps, can be used as a useful starting point for designing general purpose robotic manipulators, especially anthropomorphic ones.

5.5 Differences between Professions

The results do show significant differences in objects handled between the housekeepers and machinists. Housekeepers handle squeezable and floppy objects (such as towels and sponges) more often, whereas the machinists almost always handle rigid objects. However, in both cases cylindrical objects dominate the data set and they are virtually always grasped along the circumference (grasped dimension b/c, see Fig. 11). Also, for the other object types, the grasped dimensions are fairly stable between the professions. This suggests that the results for other professions should have many common characteristics, though future work with other professions could confirm this. Thus, the results presented should provide useful information about human object handling in general, regardless of the exact environment of interest.

6 CONCLUSIONS AND FUTURE WORK

The paper presents an object classification for grasping. The objects found both in the housekeeper and the machinist environments were classified successfully according to the presented scheme. The results can help to better understand the overall types of objects that humans frequently encounter, as well as what strategies are used to successfully grasp them.

The objects handled in our data set are surprisingly small and lightweight. In order to grasp 90 percent of the objects in the data set the way the human did, a hand should be able to grasp objects 7 cm wide and a mass of 700 g. Furthermore, it has been shown that the human has a clear tendency to grasp the smallest dimension of the object. These results can translate directly to performance specifications for a robotic hand, in terms of maximum grip aperture and payload capacity, to handle a certain percentage of common objects in human environments. Gripping by the smallest object dimension is likely a useful heuristic to use for grasp planning algorithms. Designing devices for human grasping beyond the 7 cm aperture, 700 g values may be unnecessary for human-like performance. In the environments the videos were taken, most of the objects are designed for human interaction. Therefore, the size of the actual grasp location might have been selected to fit the human hand. At this point it is unclear how those results translate to other environments that are not man-made.


In terms of future work, there are a number of limitations of the study presented here that might be addressed. While we believe that our approach provides a nice balance between effort required and generalizability, a study structure providing much more detailed measurements could help reduce uncertainty in the object properties assigned. However, we believe that the sheer number of samples considered helps to minimize any effects of imprecise object attributes. Perhaps a more important factor is related to limitations based on the nature of the grasping tasks that the subjects performed. Since we placed most of our emphasis on the correlation between grasps and object properties, we hope that the nature of the task specifics would have a minimal effect on the results, but it would still be informative to perform the study with a wider range of subjects and tasks. Overall, the current results should already be useful in defining performance specifications for robot hands, inspiring useful heuristics for grasp planners, helping to align rehabilitation goals with real-world hand usage, and to design devices for consistent grasp behavior within the comfortable range of common objects encountered.

ACKNOWLEDGMENTS

The authors would like to thank Charlotte Guertler, Joshua Zheng, and Sara De La Rosa for their work in coding the video data set used in this paper. This work was supported in part by a grant from the US National Science Foundation (NSF), grant NSF IIS-0952856.

REFERENCES

- [1] T. Iberall, "A neural network for planning hand shapes in human prehension," in *Proc. Amer. Control Conf.*, 1988, pp. 2288–2293.
- [2] T. Iberall, "Human prehension and dexterous robot hands," *Int. J. Robot. Res.*, vol. 16, no. 3, pp. 285–299, Jun. 1997.
- [3] M. R. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *IEEE Trans. Robot. Autom.*, vol. 5, no. 3, pp. 269–279, Jun. 1989.
- [4] N. Kamakura, M. Matsuo, H. Ishii, F. Mitsuboshi, and Y. Miura, "Patterns of static prehension in normal hands," *Amer. J. Occup. Therapy*, vol. 34, no. 7, pp. 437–445, Jul. 1980.
- [5] T. Feix, R. Pawlik, H.-B. Schmiedmayer, J. Romero, and D. Kragic, "A comprehensive grasp taxonomy," in *Robot., Sci. Syst.: Workshop Understanding Human Hand Advancing Robot. Manipulation*, 2009.
- [6] G. Schlesinger, *Der Mechanische Aufbau der Kunstlichen Glieder, Ersatzglieder und Arbeitshilfen*, Part II, New York, NY, USA: http://link.springer.com/chapter/10.1007%2F978-3-662-33009-8_13#, 1919.
- [7] R. L. Klatzky, B. McCloskey, S. Doherty, J. Pellegrino, and T. Smith, "Knowledge about hand shaping and knowledge about objects," *J. Motor Behav.*, vol. 19, no. 2, pp. 187–213, 1987.
- [8] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cogn. Psychol.*, vol. 19, no. 3, pp. 342–368, Jul. 1987.
- [9] L. Sartori, E. Straulino, and U. Castiello, "How objects are grasped: The interplay between affordances and end-goals," *PLoS One*, vol. 6, no. 9, p. e25203, Sep. 2011.
- [10] B. Redmond, R. Aina, T. Gorti, and B. Hannaford, "Haptic characteristics of some activities of daily living," in *Proc. IEEE Haptics Symp.*, 2010, pp. 71–76.
- [11] M. Philipose, K. Fishkin, M. Perkowitz, D. Patterson, D. Fox, H. Kautz, and D. Hähnel, "Inferring activities from interactions with objects," *IEEE Pervasive Comput.*, vol. 3, no. 4, pp. 50–57, Oct. 2004.
- [12] P. Cesari and K. M. Newell, "The scaling of human grip configurations," *J. Exp. Psychol. Hum. Perception Perform.*, vol. 25, no. 4, pp. 927–935, 1999.
- [13] P. Cesari and K. M. Newell, "Body-scaled transitions in human grip configurations," *J. Exp. Psychol. Hum. Perception Perform.*, vol. 26, no. 5, pp. 1657–1668, Oct. 2000.
- [14] R. Gilster, C. Hesse, and H. Deubel, "Contact points during multidigit grasping of geometric objects," *Exp. Brain Res.*, vol. 217, no. 1, pp. 137–151, Mar. 2012.
- [15] T. Iberall, "Grasp planning from human prehension," in *Proc. 10th Int. Joint Conf. Artif. Intell.*, 1987, vol. 2, pp. 1153–1156.
- [16] R. Balasubramanian, L. Xu, P. Brook, J. Smith, and Y. Matsuoka, "Physical human interactive guidance: Identifying grasping principles from human-planned grasps," *IEEE Trans. Robot.*, vol. 28, no. 4, pp. 899–910, Aug. 2012.
- [17] J. Z. Zheng, S. D. La Rosa, and A. M. Dollar, "An investigation of grasp type and frequency in daily household and machine shop tasks," in *Proc. IEEE Int. Conf. Robotics Autom.*, 2011, pp. 4169–4175.
- [18] I. M. Bullock, J. Z. Zheng, S. D. La Rosa, C. Guertler, and A. M. Dollar, "Grasp frequency and usage in daily household and machine shop tasks," *IEEE Trans. Haptics*, vol. 6, no. 3, pp. 296–308, Jul.–Sep. 2013.
- [19] I. M. Bullock, T. Feix, and A. M. Dollar, "Finding small, versatile sets of human grasps to span common objects," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2013, pp. 1068–1075.
- [20] I. M. Bullock, T. Feix, and A. M. Dollar, "The Yale human grasping data set: Grasp, object and task data in household and machine shop environments" (2014) [Online]. Available: <http://www.eng.yale.edu/grablab/humangrasping/>
- [21] J. Gibson, "The theory of affordances," in *Perceiving, Acting, Knowing: Toward An Ecological Psychology*, Mahwah, NJ, USA: Lawrence Erlbaum Assoc., 1977, pp. 67–82.
- [22] T. Feix, I. M. Bullock, and A. M. Dollar, "Analysis of human grasping behavior: Correlating tasks, objects and grasps," *IEEE Trans. Haptics*, 2014.
- [23] T. Zingg, "Beitrag zur Schotteranalyse," Doctoral thesis, *ETH Zürich*, Zürich, Switzerland, 1935.
- [24] D. Benn and C. Ballantyne, "The description and representation of particle shape," *Earth Surface Process. Landforms*, vol. 18, no. 7, pp. 665–672, Nov. 1993.
- [25] J. Garrett, *Clearance and Performance Values for the Bare-Handed and the Pressure-Gloved Operator*. Fairborn, OH, USA: Air Force Systems Command, Wright-Patterson Air Force Base, Aug. 1968.
- [26] R. G. Marteniuk, C. L. MacKenzie, M. Jeannerod, S. Athenes, and C. Dugas, "Constraints on human arm movement trajectories," *Can. J. Psychol.*, vol. 41, no. 3, pp. 365–378, Sep. 1987.
- [27] J. Lee Rodgers and A. Nicewander, "Thirteen ways to look at the correlation coefficient," *Amer. Stat.*, vol. 42, no. 1, pp. 59–66, Feb. 1988.
- [28] D. Moore, G. McCabe, W. Duckworth, and S. Sclove, *The Practice of Business Statistics Companion Chapter 18: Bootstrap Methods and Permutation Tests*, San Francisco, CA, USA: W. H. Freeman, 2003.
- [29] J. Hintze and R. Nelson, "Violin plots: A box plot-density trace synergism," *Amer. Statist.*, vol. 52, no. 2, pp. 181–184, May 1998.
- [30] R Core Team, "R: A language and environment for statistical computing," Vienna, Austria: R Found. Statist. Comput. 2013.
- [31] J. R. Napier, "The prehensile movements of the human hand," *J. Bone Joint Surg. (Brit)*, vol. 38-B, no. 4, pp. 902–913, Nov. 1956.
- [32] Y. Matsumoto, Y. Nishida, Y. Motomura, and Y. Okawa, "A concept of needs-oriented design and evaluation of assistive robots based on ICF," in *Proc. IEEE Int. Conf. Rehabil. Robot.*, 2011, pp. 1–6.
- [33] Y. Choi, T. Deyle, T. Chen, J. D. Glass, and C. C. Kemp, "A list of household objects for robotic retrieval prioritized by people with ALS," in *Proc. IEEE Int. Conf. Rehabil. Robot.*, 2009, pp. 510–517.
- [34] V. Mathiowetz, N. Kashman, G. Volland, K. Weber, M. Dowe, and S. Rogers, "Grip and pinch strength: Normative data for adults," *Arch. Phys. Med. Rehabil.*, vol. 66, no. 2, pp. 69–74, Feb. 1985.
- [35] R. Cuijpers, J. Smeets, and E. Brenner, "On the relation between object shape and grasping kinematics," *J. Neurophysiol.*, vol. 91, no. 6, pp. 2598–2606, Jun. 2004.
- [36] A. J. Conger, "Kappa reliabilities for continuous behaviors and events," *Educ. Psychol. Meas.*, vol. 45, no. 4, pp. 861–868, Dec. 1985.



Thomas Feix received the MSc degree in sports equipment technology from the University of Applied Sciences Technikum Wien, Vienna, Austria, and the PhD degree from the Vienna University of Technology, in 2011. He is currently a postdoctoral associate with the GRAB Lab, Department of Mechanical Engineering, Yale University, New Haven, CT. His research is focused on human grasping and manipulation and its application to robotics and prosthetics. He is a member of the IEEE.



Ian M. Bullock (S'11) received the BS degree in engineering from Harvey Mudd College, Claremont, CA in 2010. He has been working toward the PhD degree at Yale University since fall of 2010. His work experience has ranged from programming educational neuroscience modules to designing circuits for a particle accelerator. His research interests include human dexterity and robotic manipulation, with particular emphasis on using an understanding of the capabilities of the human hand to improve the design of robotic and prosthetic manipulators. He is a student member of the IEEE.



Aaron M. Dollar received the BS degree in mechanical engineering from the University of Massachusetts at Amherst, SM, and PhD degrees in engineering sciences from Harvard, and conducted two years of postdoctoral research at the MIT Media Lab. He is currently the John J. Lee associate professor of mechanical engineering and materials science at Yale. His research topics include human and robotic grasping and dexterous manipulation, mechanisms and machine design, and assistive and rehabilitation devices including upper-limb prosthetics and lower-limb orthoses. He received the 2013 DARPA Young Faculty Award, 2011 AFOSR Young Investigator Award, the 2010 Technology Review TR35 Young Innovator Award, and the 2010 US National Science Foundation (NSF) CAREER Award. He is a member of the IEEE.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**