On the Representation of Anthropomorphic Robot Hands: Shape versus Function

Ravin de Souza[*], Alexandre Bernardino[*], José Santos-Victor[*] and Aude Billard[†]

[*] Institute for Systems and Robotics
Instituto Superior Técnico
Lisboa, Portugal
{rsouza, alex, jasv}@isr.ist.utl.pt

[†] Learning Algorithms and Systems Laboratory
École Polytechnique Fédéral de Lausanne
Lausanne, Switzerland
aude.billard@epfl.ch

Abstract— **We address the problem of representations for anthropomorphic robot hands and their suitability for use in methods for learning or control. We approach hand configuration from the perspective of ultimate hand function and propose 2 parameterizations based on the ability of the hand to engage oppositional forces. These parameters can be extracted from grasp examples making them suitable for use in practical learning-from-demonstration frameworks. We propose a qualitative method to span hand functional space in a principled manner. This is used to construct a grasp set for evaluation and a qualitative baseline metric derived from human experience. Our results from human grasp data show that hand representations based on shape are not able to disambiguate hand-function. However, those based on hand-opposition primitives result in the widest separations among grasps that have radically different functions and can even clearly separate grasps whose functions overlap a great degree. We trust that these "functional parameterizations" can bridge the contrasting goals of task-oriented robotic grasping, that of controlling a dexterous robot hand to manifest hand-shape but with the ability to exercise specific hand-function.**

Keywords—hand representation; hand function; grasp synthesis; learning from demonstration

I. INTRODUCTION

The goal of any robotic grasping framework is the generation of a stable grasp that is both, feasible for a robotic hand to execute and suitable for achieving a specific task. The outcome is a hand configuration specifying the extrinsic degrees of freedom, or the position and orientation of the wrist relative to an object, and, the intrinsic degrees of freedom which make up the joint angles of the hand. The former can be thought of as the problem of "where to grasp", while the latter, a problem of "how to grasp".

Recent approaches to robot grasp synthesis adopt a data-driven or knowledge based approach, which relies on expert human knowledge about objects and how to grasp them. The main motivation is to avoid the computational complexity, and search of a large hand-object configuration space, encountered with the more traditional analytical models for grasping. For e.g. [1],[2] learn to associate successful grasps
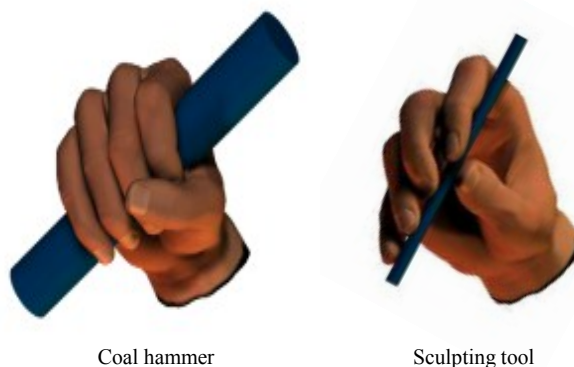

Coal hammer   Sculpting tool

Figure I-1: While the hands, in these two grasps, have similar shapes, the functions performed are very different. Pictures taken from [13]

with the local appearance of object grasping points or their relation to global shape. [3],[4] and [5] use shape decompositions of objects and principles of how to successfully grasp them, in order to generate several candidate grasp hypotheses. Alternatively, shape matching can be used to generate grasp hypotheses, either by evaluating similarities between hand-shape and object-shape [6], or by finding matching objects for which successful grasps have been previously demonstrated [7]. A recent survey by [8] mentions several other techniques for knowledge based grasping. All these techniques are oriented towards the where-to-grasp problem. In general, they report successful grasps in terms of a grasping point, an approach vector and the roll of the wrist around it. The robot hand (regardless of its flexibility) is seen as a gripper, where a grasp is completed by closing the jaws of the gripper around the object.

Anthropomorphic hands however, have the ability to organize in different shapes and forms which allow them to exert pressure along several axes simultaneously. Thus complex mixtures of power and precision are possible, which in turn can produce a rich set of effects on a grasped object. This raises the question of "how to grasp" which implies determining how to shape the hand into a prehensile posture that enables it to exert the task related forces on the object concerned and also constrains its efficacy as a sensing device able to gather information about the state of the hand-object interaction [9].

When we look at what hand representations are suitable to address these questions in the context of grasp synthesis for anthropomorphic hands, there are two principle issues of interest that arise. The first and obvious one is the simplicity (or dimensionality) of the hand representation, which directly affects the complexity of the associated grasp synthesis problem. A less obvious issue is the correspondence of hand-representation with hand-function. This aspect becomes important for supervised and semi-supervised learning approaches that attempt to capture grasping affordance relationships present in expert demonstrations of grasping. The case for "correct" representations which capture information that is actually relevant, has been well stated in [10]. In particular, with right representation, algorithms do not need to learn to ignore data and the inference problem becomes much easier to solve. For decisions related to task-oriented grasping, hand function plays a central role, and hand representations well correlated with hand function are thus better suited to learn from demonstration.

Joint angles are a preferred way of representing complex hand shapes, for the obvious reasons that they are trivial to perceive and that reconstruction of complex hand shapes is straightforward. However, joint angles are not well correlated with hand function. Santello [11], after analyzing the joint angles of 57 different grasps commonly utilized for daily tasks, found that most of these shapes could be well approximated by only 2 principle synergy directions. This suggests that the overall shape of the hand doesn't vary too much and similar shapes can be responsible for widely different functions. As can be seen in Figure I-1, a coal hammer grasp oriented towards delivering and resisting strong forces has a very similar shape to a grasp that may be employed for fine motions and moderate forces with a sculpting tool.

In this paper, we address the ability of various hand-parameterization schemes to discriminate between grasps of different function. We evaluate two kinds of computational schemes to represent the hand. "shape parameterizations" are based either on hand-shape or the mechanical structure of the hand. Joint Angles, Hand Synergies [11], Shape features [6], all fall under this category. In contrast, "functional parameterizations" are hand representations that are more oriented towards hand-function.

In section II, we provide some background on functional parameterization and motivate its ability to overcome the limitations inherent to shape parameterizations of the hand. Section III defines hand function itself and identifies a grasp data set to span this functional space in big and in small steps. A good hand representation will show adequate distance, in a consistent manner over the functional space, between points in the representation space that are clearly separated in function. We propose a baseline metric for inter-grasp distance from which a functional ordering providing expected separation between grasps can be obtained. Section IV describes in detail the various hand parameterization schemes to be compared. Section V presents the experimental setup to capture grasp data from human demonstration. Results on the correlation of the various hand-parameterization schemes with hand function are then presented in section VI and section VII concludes the paper.

II.   FUNCTIONAL PARAMETERIZATION

Functional parameterizations are hand representations more oriented toward hand-function than hand-shape.

A. Functional Categorization

There have been several attempts in the literature to classify into a set of functional categories, the myriad shapes that the human hand can achieve when grasping objects in the context of a task. The classification into precision and power grasps proposed by Napier [12], is one of the earliest and most widely cited. Since then, this broad view has been further developed into finer categories with the aim of labeling hand-shapes according to their functional abilities. More than 22 such grasp taxonomies are reviewed in [9]. Recently, a comprehensive taxonomy was provided by [13] with the goal of determining the largest set of different grasp types that can be found in the literature.

The problem with grasp taxonomies is their inability to explain all hand-shapes. We see in [14] that, "hand postures are not as discrete as most classifications suggest". In particular, any given grasp exhibits properties spanning several symbolic categories, and existing grasp classifications lack the ability to explain the myriad variations that the hand undergoes to deliver the appropriate mix of precision and power in response to the local object profile and the perceived requirements of a given task. Hence, while grasp taxonomies present an attractive one-dimensional space to describe hand-shape they are not expressive enough to explain their action (or effect causing) power.
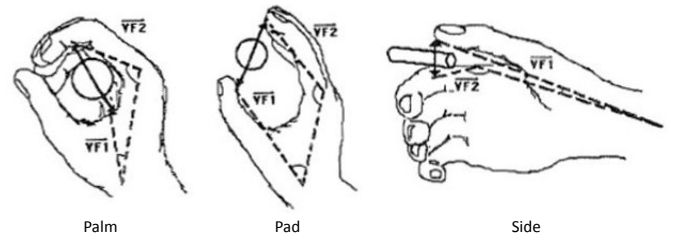
B. Opposition Space

Figure II-1. Opposition primitives. Images taken from [15]

Another way of looking at hand function computationally, is through the Opposition Space theory proposed by Iberall and Arbib [15], to explain the prehensile capabilities of the hand in a task or goal directed manner. In this theory, hand function is seen as the principal motive that drives all prehensile posturing, and the hand´s ability to engage oppositional forces are seen as a functional basis. The outwardly visible shape of the hand now becomes a side-effect of setting up a hand opposition-space in response to perceived force, motion and sensory requirements of a given task and a given object-environment context.

A functional parameterization then, attempts to parameterize an opposition space, and is guided by the following principles:

1)   *An opposition occurs between pairs of virtual fingers.* Virtual fingers were defined by Arbib [16], as parts of the hand that can work together with the same intention of opposing other hand parts or opposing task forces.

2)   *Three kinds of oppositions are possible : palm, pad, side.* These opposition primitives (Figure *II-1*) can be

combined in different ways. The precise manner in which they are combined determines the quality of the opposition space and the mix of force, motion and sensory ability that can be delivered. For a detailed explanation of opposition primitives and their properties, refer [15].

To arrive at a particular set of parameters, we must identify the oppositions that are present in a demonstrated grasp and describe the virtual fingers that form them. A satisfactory set of parameters will capture the essential qualities of the virtual fingers that contribute to hand function, and if complete, the parameter vector should be all that is needed in order to reconstruct the virtual fingers.

A functional parameterization scheme was outlined, but not implemented, by Iberall in [17]. At the tip of every virtual finger lies a *grasping surface patch:* the surface of the hand that comes in purposeful contact with the object and allows it to exercise the functions of applying oppositional pressure, imparting motion or gathering sensory information about the hand-object interaction state. Iberall in [17], proposes that a virtual finger can be satisfactorily parameterized by describing the pose and qualities of the grasping surface patch, the possible orientations of the opposition vector (or the possible force directions), the available force magnitudes for each orientation and available sensory information. However, the actual parameters used, capture a small subset of all this information. Also, it is not clear how this information relates to a practical learning from demonstration setup. In Section IV.B., we propose two parameterization schemes based on the idea that, a virtual finger's functional qualities are most highly correlated with the grasping surface patch. The first one, adapts Iberall's original scheme, introducing new parameters and measures for extracting them that are more suitable for practice. We then propose a new parameterization, significantly reducing the number of parameters required, by evaluating grasping patches in their opposition pre-shape configurations.

While shape, mechanically oriented parameterizations and grasp taxonomies, focus more on specifying and disambiguating outward shape appearance, an opposition space parameterization focusses on disambiguating functional capability at the cost of blurring (or at least not clearly specifying) the shapes that can achieve it. Our aim in this study is to find how much better (if at all) are functional parameterizations suited to discriminate between hand shapes that possess different functional abilities.

III. HAND FUNCTION SPACE

The whole problem of task-relevant grasping revolves around hand-function, representing it and matching it to task requirements given a particular environment scenario. In this study, we are interested in hand parameterizations that have a strong and consistent correlation across the hand functional space. The key questions are: What is the space of hand-function? How do you describe it? How do you measure it?

The key functions of the grasping hand can be understood along 3 dimensions: the ability to exert force/torque, the ability to impart motion/change in orientation and the ability to sense the state of the hand-object interaction [17]. Some of these dimensions can be measured analytically from a kinematic model of the hand and a description of the contacts involved in a hand-object interaction. The authors in [18],[19] list several kinds of quality metrics that can be computed. However, each metric typically focuses on a single aspect of hand function and one would need to compute an array of metrics to get a more realistic picture. Moreover these metrics assume simplified contact models and rely on precise contact information, both of which do not hold when one considers real-life demonstrations.

Another approach is to look at hand function in a more qualitative manner. Precision and power have been long been accepted as basic functional qualities by a line of researchers motivated to functionally categorize hand shape for varied reasons like prosthetics, anthropology, realistic animation and also robotic grasping and manipulation. This has resulted in a series of increasingly granular grasp taxonomies (see [9], [13] for comprehensive surveys) .

A. Grasp Data Set

Our main motivation in choosing a grasp data set is to provide a basis for evaluating the ability of various hand parameterization schemes to disambiguate hand function. For this, we adopt a qualitative approach whereby grasp function is perceived as a mix of precision and power. Accordingly, we would like to identify a representative set of grasps that spans the space of hand function, by varying precision and power abilities in big and in small steps.

Opposition space theory provides us with a principled way to do this. Essentially, palm opposition is good for providing fixing power or the general ability to resist arbitrary external wrenches. Pad opposition on the other hand provides fine dexterous ability in order to effect minute in-hand changes in object pose. Side opposition plays a supporting role to other oppositions, in some cases providing directing or aiming ability to increase resistance against task wrenches in a particular direction, or in others, providing stable grasping properties without overwhelming the ability for fine manipulation.

Using the above observations in conjunction with existing grasp taxonomies, we arrive at a suitable grasp set as follows. Note that the taxonomy of [13] was determined after examination of a vast literature on the topic and is therefore a comprehensive summary of grasps of all functions; from high power, to high precision and various combinations in between. To identify a suitable data set from this, we first manually classify these grasps into the opposition spaces to which they belong. In essence, opposition spaces are distinguished from each other by the type of oppositions employed and their mapping to real fingers. As a consequence, each opposition space offers a different mix of precision and power. Our classification differs from the one in [13] (the result of a similar procedure), because we consider all combinations of the 3 principle opposition types. We then build the data set by sampling grasps from these categories, varying the amount of power and precision delivered, and in certain cases, adding grasps when particular spaces are not well represented. Table 1 lists the final set of 17 grasps that were chosen, along with the criteria used to choose them.

B. Baseline Metric

Our strategy for evaluating how closely each hand parameterization scheme reflects hand function is based on examining whether distances between grasps, measured in

each space, bear any correlation to their separation in functional ability. For this we rely on human expert experience.

We propose a simple baseline metric for inter-grasp distance that converts the subjective qualitative assessment of a human expert on functional abilities of grasps, to a computational measure of the distance between them. The baseline is constructed by ranking all the grasps in the data set in descending order of their ability to exert power. The ranking is a qualitative assessment by the human expert. This is repeated once again for descending order of precision ability. Now, the distance between any two grasps is computed as an average of their separation in power and in precision according to their place in the respective 'expert' sorted lists. It is possible to extend this scheme by defining other qualitative functional abilities (such as ability to sense that state of hand-object interaction) and taking a weighted average of separation between grasps according to each function.

$$O_{Power} = Grasp\ set\ ordered\ by\ decreasing\ power$$

$$O_{Precision} = Grasp\ set\ ordered\ by\ decreasing\ precision$$

$$sep(O, g_1, g_2) = Ordinal\ distance\ between\ in\ g_1, g_2\ in\ O$$

$$dist(g_1, g_2) = \frac{sep(O_{Power}, g_1, g_2) + sep(O_{Precision}, g_1, g_2)}{2}$$

The scheme proposed above assumes that each qualitative function considered varies linearly across the grasp set. This is because separation is taken to be the ordinal distance in an ordered set. While the linear assumption is not necessarily true, it is still a valid assumption when using the baseline metric as a guide to find a functional ordering in the grasp set. One way of doing this to derive an ordered set of grasps of decreasing functional ability, is illustrated by Algorithm 1. Figure III-1, shows the result of applying this algorithm to the grasp data set, starting from grasp no. 10 (the highest precision grasp). The data set is therefore ordered in decreasing ability for precision.

Algorithm 1: Grasp ordering from a distance metric
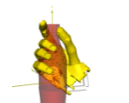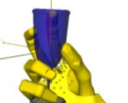
Input: Unordered set, Starting grasp (g_s), Inter-grasp distance metric $d(g_1, g_2)$

1.   Add g_s to the ordered set and delete it from the unordered set
 2.   while !empty(unordered set)
     a.   Starting grasp neighbourhood (N_{g_s}) is limited to the first 3 grasps in the ordered set.
     b.   Using $d(g_1, g_2)$, find the grasp having closest average distance to N_{g_s} from the unordered set.
     c.   Add grasp from 2.b to the ordered set and delete it from the unordered set.

Output: Ordered grasp set

Table 1. Grasp Set. Images are taken from [13] (10, 15, 16 were constructed in the GraspIt! Simulator)

|					
Strong ability to resist external wrenches from arbitrary directions.			A general ability to resist external wrenches but this is decreasing and is coupled with increasing precision ability.		
A general ability to resist external wrenches combined with increased aiming or directing ability required to keep a tool stable against particular task wrenches.		The focus here is on manipulability or dexterity.		The same focus on manipulability or dexterity, but now with increasing ability to apply stronger object torques that may be required during object use while keeping the object stably grasped. However, pad opposition is sufficient to provide both functions.	
Exactly the same objective as the previous category, but in this case the power requirements are more, necessitating the use of side opposition (in both cases by engaging thumb against the middle finger) to keep the object stably grasped while it is being manipulated.		In this case the object is big enough or the task forces involved warrant the use of palm opposition to keep the object stable while the manipulative function is being carried out. Different object or tasks mandate differing amounts of precision and power necessary.		The last category makes demands on power, dexterous ability as well as the ability to direct or aim a tool in the process of performing a task. Note that all three oppositions are employed in some mix to provide this function.	
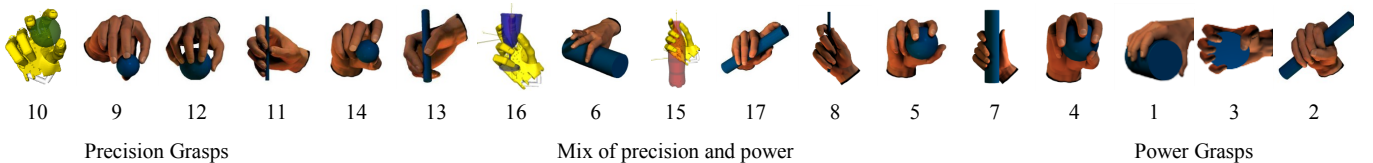
Figure III-1: Ordering of the grasp data set in decreasing order of precision obtained using Algorithm 1

IV. HAND PARAMETERIZATION APPROACHES

This section describes in more detail the various shape and functional parameter spaces we compare, and how their parameter values can be obtained from a demonstrated grasp.

A. Shape-based Parameters

1) Joint Angles: Joint angle parameters are made up of the actuated degrees of freedom of the robot hand specified in radians. For the particular hand model we use to demonstrate grasps (described in Section V), this is an 18 dimensional vector.

2) Hand Synergies: A synergy parameterization of a demonstrated grasp is a 6 dimensional vector obtained by projecting the joint angle vector onto a lower dimensional subspace formed by hand synergies. Following the procedure described by Santello in [11], principal component analysis is applied to a set of 57 commonly used grasps, and a sub-subspace accounting for 90% of the variance in the joint angle data is identified. 6 principal components are required for this. Non-linear techniques for sub-space reduction could further improve these results.

For the case of joint angles and synergies, a hand configuration is a point in $\mathbb{R}^d$. Hence, the distance between two points is taken as the $l_2 - norm$ of their vector difference

3) Shape Features: According to Li et al [6], the shape characteristics of each demonstrated grasp can be captured by a compound feature consisting of a set of 3 dimensional features obtained from a point cloud. The point clouds come from a descretization of the grasping surface patches present in a demonstrated grasp. Following the method outlined by the authors, distance between two demonstrated hand-shapes, A and B, is a weighted average of the features in the feature set for grasp A, to their nearest neighbours in the feature set for grasp B. The weight of a grasp feature, determines its importance to the overall shape of the grasp in which it occurs, and is proportional to its occurrence in that grasp with respect to its occurrence in all other grasps in the data set. An outcome of this definition is that the weights of grasp A features and grasp B features will be different (as they are different shapes) and hence the distance measure will not be symmetric, i.e. $dist(A, B) \neq dist(B, A)$. To overcome this, we take the distance between two grasps as the average of the distances computed for both directions.

B. Functional Parameters

All the opposition space parameterization schemes are defined in a hand-centric space. As depicted in Figure IV-1, this is a 3D co-ordinate frame centered in the palm at the wrist. A hand-shape is thus enclosed in a box where the palm surface lies on the X-Y plane and the height along the Z-axis is determined by the reach of the fingers above the palm. All demonstrated grasps are first transformed to this reference frame before parameters describing the virtual fingers are extracted. Furthermore, all angle parameters are multiplied by a factor of $1\ cm/rad$ to bring them on par with distance parameters. Each point in the opposition space

parameterizations is a vector in $\mathbb{R}^d$ and distance is measured as the $l_2 - norm$ of the difference between two points.

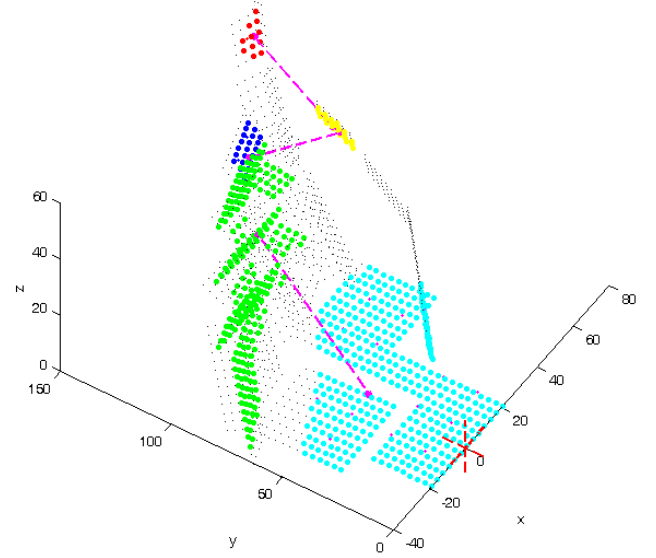*4) VirtualFinger1 – based on the demonstrated grasp*

Figure IV-1: Shows data for grasp no. 17 in the 3D hand centric frame. A virtual finger is characterized a grasping patch, and an opposition, by a pair of virtual fingers. Grasping patch point clouds (points in color) and opposition vectors (lines in magenta) of the demonstrated grasp are overlaid on the grasping surfaces of the hand (points in black). Best viewed in color.

This parameterization describes the virtual finger pairs comprising the oppositions present in a demonstrated grasp. Using the model presented earlier (section II.B.), a virtual finger is characterized by a grasping patch. As seen in Figure IV-1, for each demonstrated grasp, the set of 3D point clouds and opposition vectors associated with opposing grasping patches, is available to us. From these can be extracted 8 scalar parameters per virtual finger as listed below. One group of parameters describes the focus of oppositional pressure, and another group describes the grasping surface patch itself.

p_1, p_2, p_3	Coordinates (x,y,z) of the focus of opposition
p_4, p_5	Each grasping patch is approximated by a plane defined by the directions of maximum and minimum variance of the corresponding grasping patch point cloud. p_4 and p_5 constitute bounds of the point cloud projected onto this plane.
p_6, p_7, p_8	p_6 and p_7 refer to the azimuth and elevation of the grasping patch plane normal, and p_8 refers to the roll of the plane with respect to the opposition vector

Note that, this set of 8 parameters describes 1 virtual finger. To complete the hand-parameterization, it is necessary to describe 2 virtual fingers for each opposition: palm, pad and side. This makes for a total of 48 parameters. If an opposition is not present in the demonstrated grasp, its virtual fingers are described by the zero vector.

5) VirtualFinger2 – based on opposition pre-shape

Palm pre-shape

Pad pre-shape

Side pre-shape

VF1 $[p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$ $\quad [p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$ $\quad [p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$

VF2 $[p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$ $\quad [p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$ $\quad [p_1 \quad p_2 \quad p_3 \quad p_4 \quad p_5 \quad p_6 \quad p_7 \quad p_8]$
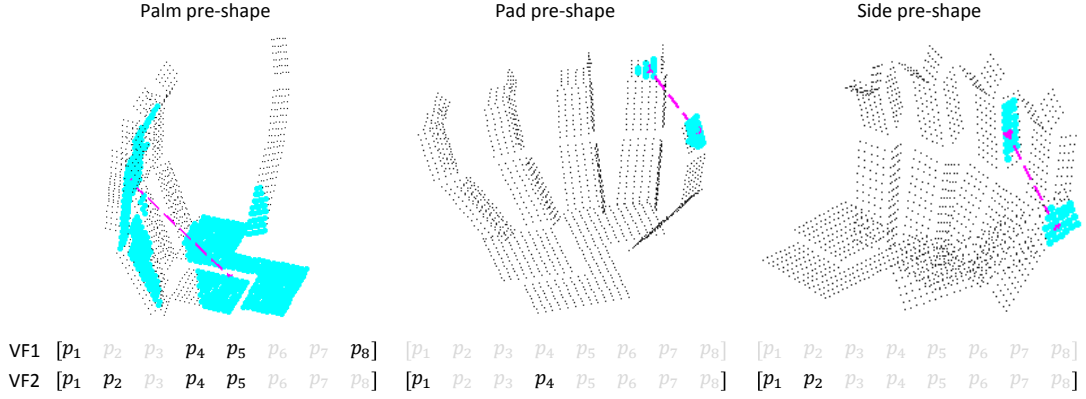
Figure IV-2: *VirtualFinger2* parameterization scheme. The figure shows the data for grasp no. 17. However in contrast to Figure IV-1, grasping patches and opposition vectors are now examined in the pre-shape configuration associated with each opposition type. *VirtualFinger2* parameters are listed below each pre-shape. They have the same meanings as in *VirtualFinger1,* however, several parameters can be fixed to their pre-shape configuration values and have been grayed out. Best viewed in color.

Our main proposal is a novel parameterization approach which looks at the prototype grasp (or grasps) from which the final demonstrated grasp could have originated. Consider that there exists an opposition pre-shape for each opposition type. This is a known fixed configuration of the hand, which captures the intention of a particular opposition type and which serves as a starting point for hand closure, in order to arrive at a final grasp manifesting that opposition. Working from this assumption, the oppositions detected in a demonstrated grasp, are the outcome of their respective opposition pre-shapes being mixed together by the closure of their respective grasping patches along their respective opposition vectors (the process referred to as enclose-schemas in [15]). The closure of the hand to form the final shape always operates under the basic functional intention determined by the set of pre-shapes from which it originated, and hence it is sufficient to examine these, instead of the final grasp. As closure progresses, it becomes subject to other influences more concerned with satisfying hand kinematic constraints and ensuring compliance of the hand with the object surface in order to form a stable grasp.

Figure IV-2 describes the *VirtualFinger2* set of 12 parameters. These parameters are exactly the same as described earlier for *VirtualFinger1,* as they arise from the same virtual finger model. However, due to the constraint of grasping patches being bound to opposition pre-shape configurations, full flexibility in 3D hand space is not required and several parameters can be fixed to their pre-shape values. For instance, in the example of Figure IV-2, we describe pad opposition using only the finger pad grasping patch, and within this, only the x-coordinate of opposition focus and grasping patch length along direction of maximum variance.

V. EXPERIMENTAL SETUP

An experimental framework was constructed wherein human demonstration grasp data could be collected for evaluation. The human experimenter demonstrates grasps using a Cyberglove sensor to control the 18-dof Shadow Robot hand model in the GraspIt! Simulator [20]. The experimenter first demonstrates a sequence of known hand-shapes for the purpose of calibration. This allows the sensors of the Cyberglove to be mapped to the joint angles of the robot hand model using linear regression. Once the ability to control the model is validated, the experimenter demonstrates each of the grasps listed in Table 1, using the same objects. For each grasp that is demonstrated, we collect information required to extract hand parameter values. This consists of:

1)   Joint Angles. Obtained directly from the Cyberglove sensors after suitable transformation.

2)   Grasping Patches and Location of Opposition Foci within them. Obtained through a process of manual annotation. The experimenter is presented with a flat (2D) view of the grasping surfaces of the hand. He/she then proceeds to delimit grasping patch polygons and the opposition vectors on it. This information is then transformed to the 3D hand centric frame using joint angles and known forward kinematics. This provides us with a point cloud for each grasping patch and a vector in 3D space for each opposition.

Note that, the process of delimiting grasping patches is far from precise. It is based on visual estimation in the mind of the human annotator and hence, is subject to variation and noise. The same information could also be obtained by employing an array of tactile sensors on the grasping surfaces of the demonstrating hand (simulated or real). Such a capability would be necessary in a practical learning from demonstration setup. However, for the purposes of this study, manual annotation is sufficient.

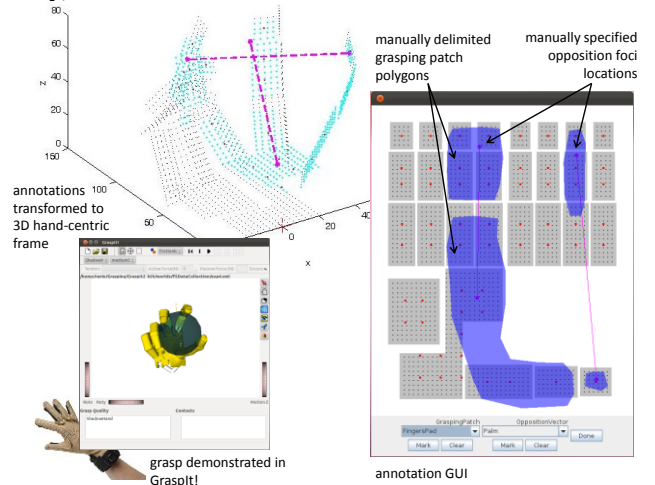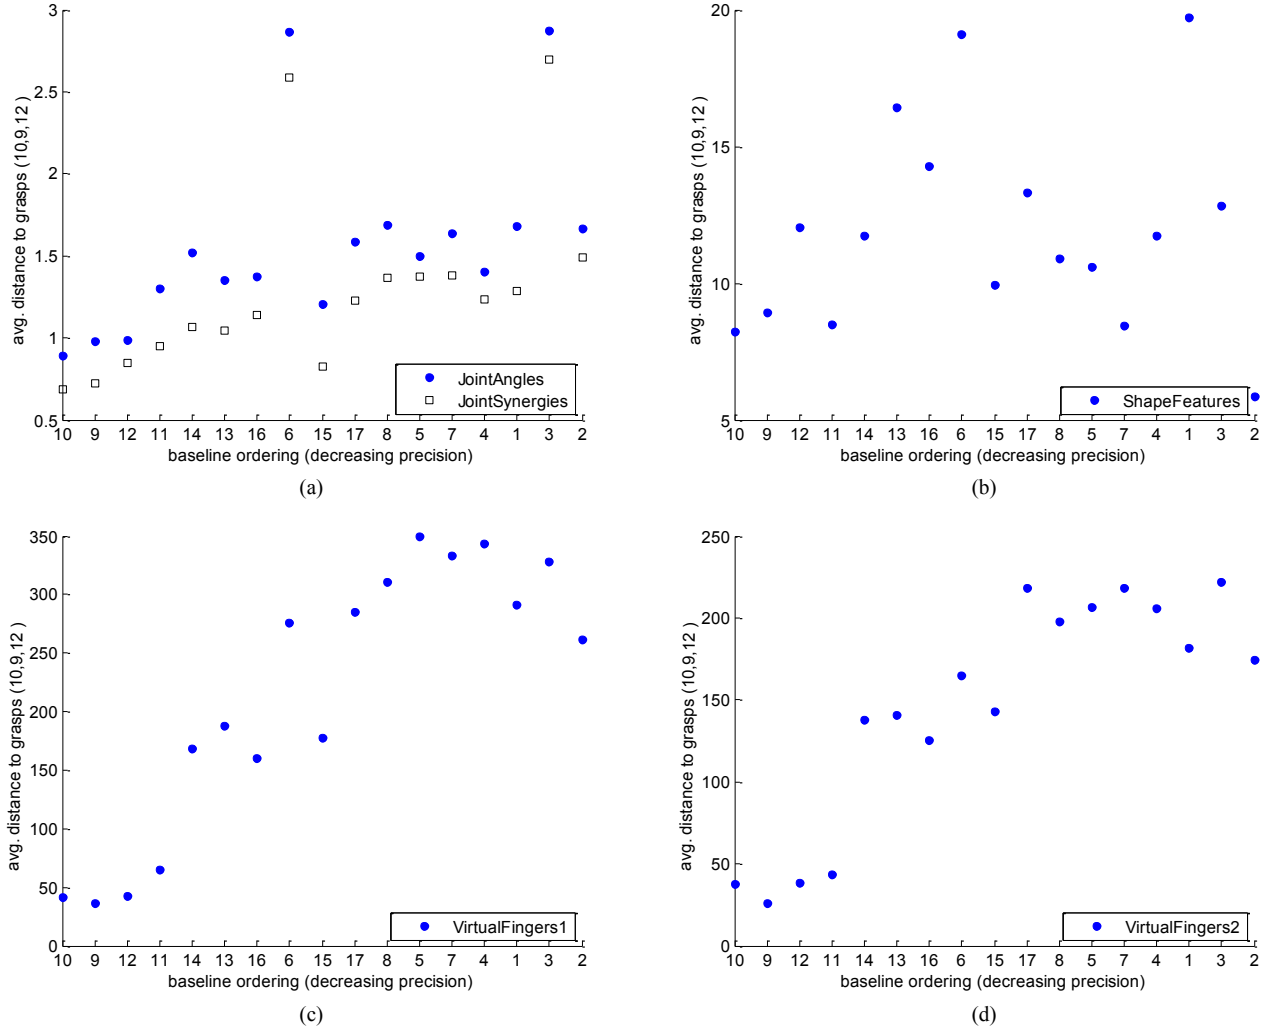annotations transformed to 3D hand-centric frame

manually delimited grasping patch polygons

manually specified opposition foci locations

grasp demonstrated in GraspIt!

annotation GUI

Figure V-1: Framework for grasp data collection. Best viewed in color.

(a)

(b)

(c)

(d)

	JointAngles	*JointSynergies*	*ShapeFeatures*	*VirtualFinger1*	*VirtualFinger2*
Pearson Coefficient	0.5642	0.5885	0.0803	0.8629	0.8376
p-Value	0.0183	0.0129	0.7592	0	0

(e)

Figure V-2: Examines how well inter-grasp distances obtained from the various parameterization schemes compare with functional orderings of the grasp set based on human expert experience (and the inter-grasp separation thereby induced). (a), (b), (c) and (d) plot distance to highest precision for the grasp set ordered by the baseline metric in decreasing precision ability (see Figure III-1). Table (e) shows the correlation of each plot to the same distances obtained from the baseline metric. p-values indicate the probability that a particular correlation occurs by chance.

VI. RESULTS

In this section we examine the ability of each parameterization scheme to disambiguate hand function, by comparing inter-grasp distances against separations mandated by the baseline ordering. As shown in Section III.B, the human experience baseline metric can be used to obtain orderings of the grasp data set based on hand function. In particular, we make use of the ordering in Figure III-1-a: *decreasing precision ability*. This ordering induces a separation among grasps in the data set that must also be respected by hand-parameterizations that are correlated with hand function.

Figure V-2 reports distances obtained from the various parameterization schemes. In each case, the x-axis is the human baseline ordering of decreasing precision ability. The y-axis plots the average distance of each grasp to a small neighbourhood, $N = \{10,9,12\}$, of the highest precision grasp. As we have taken care to choose our data set to adequately span the space of precision and power, we expect to see an increasing trend (not necessarily linear), where distance from N increases for grasps further away from N in the baseline ordering.

Referring to Figure V-2-{a,b}, the distances reported by the shape based parameterizations display no clear separation between precision and power. With *JointAngles* and *JointSynergies*, a slight increasing trend can indeed be noticed, however, for the majority of the data set, distance to N varies in a narrow band (0.6-1.6). The data for *ShapeFeatures* is erratic and doesn´t show any discernible

trend at all. We also notice that the shape based parameterizations schemes show anomalies, where grasps clearly strong in power are reported closer to the high-precision neighbourhood than other more precision oriented grasps. Anomalies are more pronounced in the case of *ShapeFeatures* (consider for example grasps {7,11} in Figure V-2-b).

In contrast to this, Opposition Space parameterizations are clearly able to distinguish power from precision. Interestingly, we see emerge from the distance data, 3 categories of grasps that agree closely with the inter-grasp separations mandated by the human experience baseline. For instance, grasp {11} is close to N in the ordering and also in distance. Similarly grasps {4,7} which are power oriented and, are positioned at the far end of the baseline ordering, are also widely separated from N in distance measured in parameter space. Finally, grasps {15,16} close to the middle of the ordering, end up positioned between the strong precision and strong power categories in distance to N.

This correlation with human intuition can be expressed numerically by computing the Pearson correlation coefficient for each parameterization scheme against the human experience baseline (Figure V-2-e). This is done by dividing, in each case, the covariance of two distance sets, one from parameter space and the other from the baseline metric, by the standard deviations of each set. The high correlation coefficients associated with the opposition space parameterizations suggest that these schemes are better than others in mimicking the thought process of the human expert, to discriminate grasps close and far in precision and power.

Two further observations are also in order. First, the range of distance data for the opposition space parameters (0-350 and 0-250), is much larger than that of the shape based ones (0-3 and 0-20). This indicates that opposition space parameters are more sensitive to variance in function. Second, distance data from the scheme based on pre-shape, *VirtualFinger2*, is very similar to scheme based on demonstrated configuration, *VirtualFinger1*. This indicates that full flexibility of the grasping surface patch (and the significant increase in number of parameters) doesn´t add much more in disambiguating hand-function.

VII. Conclusions

In this paper we have compared shape based representations of the hand to those based on the ability of the hand to engage oppositional forces. We proposed 2 parameterizations of opposition space based on virtual fingers that model grasping patches and the opposition foci within them. A novel approach that examines grasping patches in their pre-shape configurations yielded a significant reduction in the number of parameters required. We proposed a qualitative method to span hand functional space in a principled manner. This was subsequently utilized to select a set of grasps and define baseline metric over which the correlation of hand representation with hand function could be evaluated.

We find that parameterizations based on opposition primitives exhibit a strong correlation with hand function and are able to mimic functional orderings and categorizations of grasps as suggested by human experience. In addition, these computational schemes to represent the hand, are consistent

over hand function space and sensitive to small as well as big changes in power and precision. These properties are not observed in the several shape parameterizations schemes also evaluated. Moreover, we find that imposing opposition pre-shape constraints on the virtual fingers, does not affect their correlation with hand-function.

In future work we look to use these "functional parameters" along with probabilistic techniques to learn associations between hand-object interaction and the resulting changes that can be effected on a grasped object. Such information can help answer questions on "how-to-grasp" and also provide essential data, in terms of a dynamic model, for planning in-hand manipulations.

References

[1] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic Grasping of Novel Objects using Vision," *The Int. Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.

[2] J. Bohg and D. Kragic, "Learning grasping points with shape context," *Robotics and Autonomous Systems*, vol. 58, no. 4, pp. 362–377, 2010.

[3] K. Huebner and D. Kragic, "Selection of robot pre-grasps using box-based shape approximation," in *IEEE Int. Conf. on Intelligent Robots and Systems*, 2008, pp. 1765–1770.

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

[5] C. Goldfeder, P. K. Allen, C. Lackner, and R. Pelossof, "Grasp Planning via Decomposition Trees," in *IEEE International Conference on Robotics and Automation*, 2007, no. April, pp. 4679–4684.

[6] Y. Li, J. L. Fu, and N. S. Pollard, "Data-driven grasp synthesis using shape matching and task-based pruning.," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 4, pp. 732–747, 2007.

[7] C. Goldfeder and P. K. Allen, "Data-driven grasping," *Autonomous Robots*, vol. 31, no. 1, pp. 1–20, 2011.

[8] A. Sahbani, S. El-Khoury, and P. Bidaud, "An overview of 3D object grasp synthesis algorithms," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 326–336, 2011.

[9] T. Iberall, "Human prehension and dexterous robot hands," *The International Journal of Robotics Research*, vol. 16, no. 3, pp. 285–299, 1997.

[10] C. H. Ek and D. Kragic, "The importance of structure," in *International Symposium on Robotics Research*, 2011.

[11] M. Santello, M. Flanders, and J. F. Soechting, "Postural hand synergies for tool use.," *Journal of Neuroscience*, vol. 18, no. 23, pp. 10105–10115, 1998.

[12] J. R. Napier, "The prehensile movements of the human hand," *Surgery*, vol. 38-B, no. 4, pp. 902–13, 1956.

[13] T. Feix, R. Pawlik, H. 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.

[14] T. Iberall, "The nature of human prehension: three dexterous hands in one," in *IEEE Int. Conf. on Robotics and Automation*, 1987, vol. 4, pp. 396–401.

[15] T. Iberall, G. Bingham, and M. A. Arbib, "Opposition space as a structuring concept for the analysis of skilled hand movements," *Experimental Brain Research*, vol. 15, pp. 158–173, 1986.

[16] M. A. Arbib, T. Iberall, and D. Lyons, "Coordinated control programs for movements of the hand," *Experimental Brain Research*, vol. 10, pp. 111–129, 1985.

[17] T. Iberall, C. Torras, and C. Mackenzie, "Parameterizing Prehension: A mathematical model of opposition space," in *Cognitiva*, 1990.

[18] 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.

[19] R. Suárez, M. Roa, J. Cornella, and J. Cornell, "Grasp quality measures," *Control*, 2006.

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