UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

LEARNING GRASP AFFORDANCES

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By

CHARLES DE GRANVILLE
Norman, Oklahoma
2008

LEARNING GRASP AFFORDANCES

A THESIS APPROVED FOR THE
SCHOOL OF COMPUTER SCIENCE

BY

Dr. Andrew H. Fagg

Dr. Amy McGovern

Dr. Dean Hougen

Dr. Robert Platt Jr.

Dedication

This thesis is dedicated to Chellee Semon for putting up with me all of these years.

Acknowledgments

This research was supported by NSF/CISE/REU award #0453545 and by NASA EPSCoR grant #7800. I would like to thank Joshua Southerland, Di Wang, and Brian Watson for all of their help and support. In addition, I would also like to thank Rob Platt, Mike Goza, Myron A. Diftler, and William Bluethmann at the Dexterous Robotics Laboratory at NASA's Johnson Space Center for their support in providing access to Robonaut. Finally, I would like to give special thanks to Dr. Andrew Fagg for guiding and supporting me throughout the course of this research.

Contents

List Of Figures

Abstract

When presented with an object to be manipulated, a robot must identify the available forms of interaction. How might an agent acquire this mapping from object representation to action? In this thesis, I describe an approach that learns a mapping from objects to grasps through human demonstration and robot teleoperation. For a given object, the teacher demonstrates a set of feasible grasps. These example grasps are then clustered in terms of the pose of the hand and the configuration of the fingers. Individual clusters in this space are represented using probability density functions, and thus correspond to variations around canonical grasp types. Multiple clusters are captured through a mixture distribution-based representation. Experimental results demonstrate the feasibility of extracting a compact set of canonical grasps from the provided experience. Each of these canonical grasps can then be used to parameterize a reach controller that brings a robot hand into a specific spatial relationship with an object.

Chapter 1

Introduction

When designing machines to autonomously operate in complex environments, many interesting challenges arise. For example, suppose a robot is tasked with building a structure. The robot is endowed with a visual system, arms, hands, and the ability to locomote. What specific problems need to be solved in order for it to successfully complete the task? First, the robot must be able to recognize and locate the objects in its environment. Second, it may also need to navigate to a desired location within the workspace, avoiding obstacles if necessary. An important component of any construction task is manipulation. Hence, the ability to pick up objects and place them in desired configurations is crucial. Finally, the robot must devise a control policy that describes the sequence of actions necessary to build the structure. This thesis focuses on the third step in this process: the planning and execution of grasping actions.

Grasping, however, proves to be a difficult challenge in the robotics domain. Even if the robot is able to recognize the identity and pose of an object, there may be a large number of hand configurations that lead to a successful grasp. From the perspective of the agent, it is not obvious how to derive these solutions *a priori*. Furthermore, a large number of grasping options complicates higher level planning and learning. Thus, a means of compactly describing the viable grasping options for an object is needed.

Gibson (1966, 1977) suggested that objects in the environment can be represented by an agent in terms of the actions that can be made with respect to those objects. Furthermore, he suggested that these representations should be distinct from those that explicitly capture the physical properties or semantics of the objects. This *affordance* representation captures the combination of the *interaction-relevant* physical

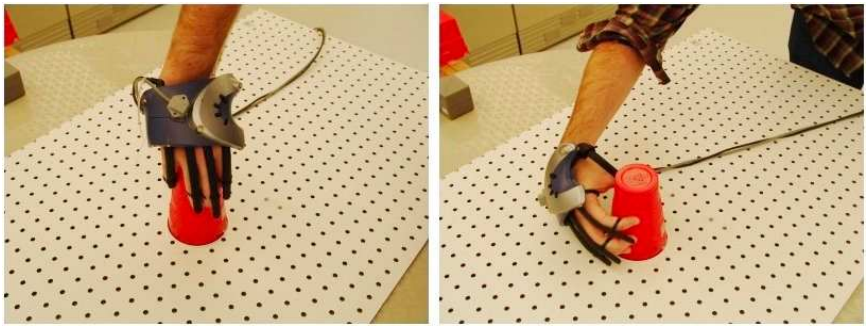

Figure 1.1: Cup grasps.

properties of the object and the capabilities of the agent's body. Object properties that are not relevant to action are implicitly lost. An affordance is thus an abstraction mechanism that specifies what actions are possible with an object. For example, when grasping the cup shown in figure 1.1 there are essentially two options; an approach from the top (a), and an approach from the side (b). Furthermore, the hand can rotate arbitrarily about the vertical axis through the cup's center when using either option. This description is compact, yet it captures a large number of the possible grasps that may be made with respect to the cup.

I hypothesize that affordances are an effective means for robots to represent and reason about the world around them. In this thesis, I propose an approach in which a robot *learns* the grasp affordances of objects through human demonstration and robot teleoperation. More specifically, for a given object the human teacher demonstrates a set of feasible grasps. These example grasps are then clustered in terms of the pose of the hand and the configuration of the fingers. Individual clusters in this space are represented using probability density functions, and thus correspond to variations around canonical grasp approaches. Multiple clusters are captured through a mixture distribution-based representation. I hypothesize that this clustering approach is an effective means for constructing a compact representation of how an object may be grasped. Furthermore, I hypothesize that the learned clusters can then be used by a robot to plan and execute grasping actions.

This thesis is organized as follows: Chapter 2 reviews related work on robotic grasping, Chapter 3 describes probability density functions over hand pose, Chapter 4 details a method for learning the canonical hand poses used to grasp an object based on a human demonstration, Chapter 5 applies the methods of Chapters 3 and 4 to

experience generated through robot teleoperation, Chapter 6 introduces an approach to modeling hand pose and finger configuration, Chapter 7 proposes an algorithm that uses the learned affordance representation for planning and executing grasps, and Chapter 8 evaluates this algorithm experimentally using a real robot.

Chapter 2

Related Work

One important form of interaction is that of grasping. For a given object, how might an agent come to represent the set of feasible grasps that may be made? Ultimately, one must establish a mapping from perceivable features to a set of parameterized grasping actions (specific positions and orientations for the hand, as well as configurations for the fingers) that are expected to be successful if executed. We would like for this set to be small so as to facilitate both planning and exploratory learning. One approach is to establish a *direct* mapping between the perceived features of objects and the grasps available to the agent.

Coelho, Piater, and Grupen developed an approach that automatically learns a direct mapping from constellations of visual features to hand orientations and finger configurations in a planar grasping task (Coelho et al. 2000; Piater and Grupen 2002). A haptically-driven approach is used to explore a given object in order to find a set of finger contact locations that minimize the net force and torque exerted upon the object, which results in a collection of stable grasp configurations (Coelho and Grupen 1997). Primitive image features consist of directed edges and textures, while constellations of features define the geometric relationships between multiple primitives. This robustly captures shape over arbitrary areas of an image. Given a set of object images and the identified stable grasp configurations, the visual learning algorithm attempts to find constellations that consistently predict the relative orientation of the hand and the finger configuration. A candidate visual constellation is evaluated with respect to the set of example grasps by identifying clusters of relative hand orientations. When a small number of dense clusters has been identified, the

visual constellation is considered to be predictive of hand orientation. In novel situations, these affordance maps can then be used to position and configure the hand in such a way that a successful grasp is probable.

Similarly, Sweeney and Grupen (2007) present an approach that learns a collection of shared grasp affordances for a set of objects. Each affordance is represented as a joint probability distribution over coarse visual features, hand position, and hand orientation. Visual features are computed by first eliminating the pixels that belong to the image background, and then fitting an ellipse to the remaining object pixels. The features of the ellipse, such as its center and the orientation of its axes, serve as a rough estimate of the object's pose in the image. Successful examples of grasps are demonstrated by a human teleoperator for a variety of objects, and the provided experience is described in terms of a small number of clusters. Each cluster represents an affordance, and specifies a mapping from similar visual features to specific positions and orientations for the hand. When the visual features exhibited by a novel object are similar to those in the training set, the learned affordance map can be used to derive actions that are likely to result in a successful grasp.

Montesano et al. (2007) introduce an approach that establishes relationships between object features, actions, and their corresponding effects. These relationships are represented by a graphical structure called a Bayesian network, and are learned through experience autonomously generated by a humanoid robot. Nodes in the graph represent object features, actions, and effects, while directed edges correspond to causal relationships between nodes. The structure of the graph represents a probability distribution that can be used for prediction, planning, and action recognition. As an example, suppose the robot taps a ball with its hand, which causes the ball to roll. In this scenario, the graph structure learned by the algorithm assigns a high probability to the outcome of large object velocity given that the shape is spherical and the object is tapped. When subsequently presented with an object that is spherical, the robot is able to predict that it will move quickly if tapped. Conversely, if the object is square, the probability of observing high object velocity is low. Furthermore, information that is irrelevant to prediction is implicitly lost (e.g., the color of the ball does not affect whether or not it will roll if it is tapped).

One challenge that arises out of taking a direct approach to learning grasp affordances is that the appearance of an object may change as a function of its pose even

though the set of grasps that may be made relative to the object does not. One possible solution to this problem is to establish an *indirect* mapping from sensory features to grasps through an intermediate object-centered representation. This approach separates the problem into one of first estimating an object's pose and identity (and/or shape), and then estimating the set of feasible configurations of the hand relative to the object. Stoytchev (2005) presents a developmental approach to the latter of these two problems. This approach discovers sequences of actions that result in a successful "binding" of the object with the robot (i.e., a grasp). The robot is initially endowed with a set of actions (involving movement of the arm, as well as opening and closing of the hand) and a set of visual operators that are relativized to various points attached to the robot. However, the robot has no explicit representation of hands or grasping. For each of several objects, the robot performs a random sequence of exploratory actions. Subsequences of actions that lead to simultaneous movement of the object and a component of the robot are deemed as "interesting." Short subsequences that reliably achieve this interesting *bound* configuration are identified as viable grasping actions, and are associated with the object through the affordance map. Hence, when a known object is subsequently presented to the robot, it is able to generate a sequence of actions that will likely lead to a successful grasp.

In contrast to making use of object identity, Bekey et al. (1993) and Miller et al. (2003) rely on a shape-based description of an object. Objects are modeled as a collection of shape primitives such as cylinders, rectangular prisms and cones. Each primitive is associated with a set of relative hand poses and corresponding finger configurations. For example, the palm positions for a rectangular prism consist of a finite number of points along each of its faces. The distance of each point from a face, as well as the number of points along a face, are determined by the size of the rectangular prism. For each palm position, the corresponding palm orientation is parallel to the nearest surface. The finger configurations used for each grasp point ensure that the thumb and remaining fingers contact opposing faces of the rectangular prism. Given a set of primitive shapes that describe a novel object, the planner can quickly generate a set of candidate grasps. In Miller's case, each candidate grasp is executed in simulation, and evaluated in terms of the amount of force required to dislodge the object from the hand. In the case of Bekey et al., the set of candidates is evaluated as a function of the semantics of the grasp and task. For example,

when using a wrench to turn a nut, grasps that provide a large amount of torque are preferable.

Chapter 3

Probability Density Functions over Hand Pose

3.1 Representing Grasp Affordances

Having a compact representation that describes the ways in which an object may be grasped greatly improves an agent's ability to efficiently plan and execute grasping actions. Our goal is to compress a large number of examples provided by a human teacher into a small number of clusters that are meaningful in terms of describing the functionally different ways that an object may be grasped. Our approach represents each of these clusters as a probability density function (PDF) defined over both orientation and position. A set of clusters is then captured using a mixture model approach.

3.2 PDF's in Orientation Space

Unit quaternions are a natural representation of 3D orientation because they comprise a proper metric space, a property that allows us to compute measures of similarity between pairs of orientations. Here, an orientation is represented as a point on the surface of a 4D unit hypersphere. This representation is also antipodally symmetric: pairs of points that fall on opposite poles represent the same 3D orientation. The Dimroth-Watson distribution captures a Gaussian-like shape on the unit hypersphere, while explicitly acknowledging this symmetry (Mardia and Jupp 1999; Rancourt et al. 2000). The probability density function for this distribution is as follows:

$$f\left(\mathbf{q}|\mathbf{u},k\right) = F\left(k\right) e^{k\left(\mathbf{q}^T\mathbf{u}\right)^2}, \tag{3.1}$$

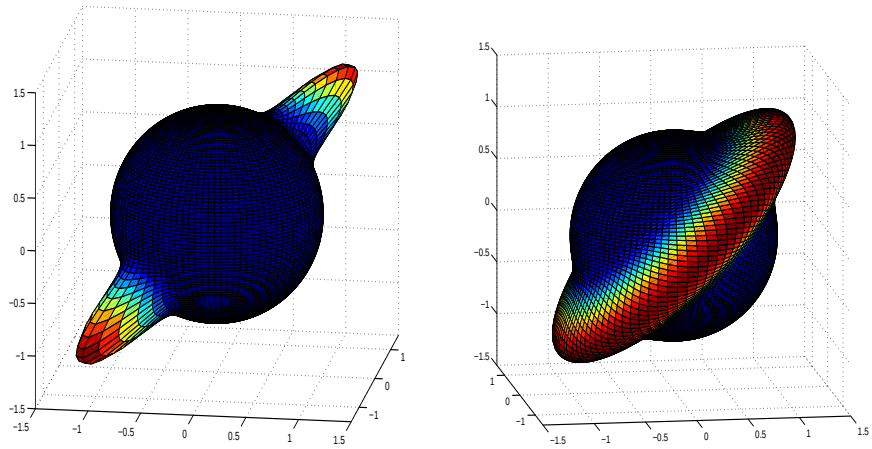

(a) (b)

Figure 3.1: Three dimensional representations of the Dimroth-Watson (a) and girdle (b) distributions on S2. In both cases, the surface radius is $1 + p$, where p is the probability density at the corresponding orientation

where $\mathbf{q} \in \mathbb{R}^4$ is a unit quaternion, $\mathbf{u} \in \mathbb{R}^4$ is a unit quaternion that represents the "mean" rotation, $k \geq 0$ is a concentration parameter, and $F(k)$ is a normalization term that is derived in the appendix. Note that $\mathbf{q}^T\mathbf{u} = \cos\theta$, where θ is the angle between $\mathbf{q}$ and $\mathbf{u}$. Hence, density is maximal when $\mathbf{q}$ and $\mathbf{u}$ are aligned, and decreases exponentially as $\cos\theta$ decreases. When $k = 0$, the distribution is uniform across all rotations; as k increases, the distribution concentrates about $\mathbf{u}$. Figure 3.1(a) shows a 3D visualization of the Dimroth-Watson distribution, and highlights its Gaussian-like characteristics. The high density peaks correspond to $\mathbf{u}$ and $-\mathbf{u}$.

A second cluster type of interest corresponds to the case in which an object exhibits a rotational symmetry. For example, an object such as a cylinder can be approached from any orientation in which the palm of the hand is parallel to the planar face of the cylinder. In this case, hand orientation is constrained in two dimensions, but the third is unconstrained. This set of hand orientations corresponds to an arbitrary rotation about a fixed axis, and is described by a great circle (or girdle) on the 4D hypersphere. We model this set using a generalization of the Dimroth-Watson distribution that was suggested by Rivest (2001). The probability density function is as follows:

$$\bar{f}\left(\mathbf{q}|\mathbf{u}_1, \mathbf{u}_2, k\right) = \bar{F}\left(k\right) e^{k\left[\left(\mathbf{q}^T\mathbf{u}_1\right)^2 + \left(\mathbf{q}^T\mathbf{u}_2\right)^2\right]}, \tag{3.2}$$

where $\mathbf{u}_1 \in \mathbb{R}^4$ and $\mathbf{u}_2 \in \mathbb{R}^4$ are orthogonal unit quaternions that determine the great circle, and $\bar{F}(k)$ is the corresponding normalization term that is derived in the appendix. Figure 3.1(b) illustrates the girdle distribution on S2 (the 3D sphere). First, note that all points on the great circle are assigned maximal density. This corresponds to the set of points for which $\left(\mathbf{q}^T\mathbf{u}_1\right)^2 + \left(\mathbf{q}^T\mathbf{u}_2\right)^2 = 1$. However, also notice that as the cosine of the angle between $\mathbf{q}$ and the point closest to $\mathbf{q}$ on great circle decreases, density decreases exponentially.

For a given set of observations, the parameters of the Dimroth-Watson and girdle distributions are estimated using maximum likelihood estimation (MLE). The axes of the distribution are derived from the sample covariance matrix, $\mathbf{\Lambda} \in \mathbb{R}^{4\times 4}$:

$$\mathbf{\Lambda} = \frac{\sum_{i=1}^{N} \mathbf{q}_i \mathbf{q}_i^T}{N}, \tag{3.3}$$

where $\mathbf{q}_i$ is the orientation of the ith sample, and N is the total number of samples. The MLE of $\mathbf{u}$ is parallel to the first eigenvector of $\mathbf{\Lambda}$ (Mardia and Jupp 1999; Rancourt et al. 2000). The orthogonal vectors $\mathbf{u}_1$ and $\mathbf{u}_2$ span the same space as the first and second eigenvectors of $\mathbf{\Lambda}$ (Rivest 2001).

For the Dimroth-Watson distribution, the MLE of the concentration parameter, k, uniquely satisfies the following (see the appendix for the derivation):

$$G\left(k\right) \equiv \frac{F'\left(k\right)}{F\left(k\right)} = -\frac{\sum_{i=1}^{N}\left(\mathbf{q}_i^T\mathbf{u}\right)^2}{N}. \tag{3.4}$$

In the case of the girdle distribution, the MLE of k uniquely satisfies (see the appendix for the derivation):

$$\bar{G}\left(k\right) \equiv \frac{\bar{F}'\left(k\right)}{\bar{F}\left(k\right)} = -\frac{\sum_{i=1}^{N}\left[\left(\mathbf{q}_i^T\mathbf{u}_1\right)^2 + \left(\mathbf{q}_i^T\mathbf{u}_2\right)^2\right]}{N}. \tag{3.5}$$

For computational efficiency, we approximate $G^{-1}()$ and $\bar{G}^{-1}()$ when solving for k (see the appendix for details).

3.3 PDF's in Cartesian Space

The position of the hand is represented as a 3D vector in Cartesian space. We choose to model position using a Gaussian distribution:

$$p\left(\mathbf{x}|\boldsymbol{\mu}, \mathbf{\Sigma}\right) = \frac{1}{\left(2\pi\right)^{\frac{d}{2}}\left|\mathbf{\Sigma}\right|^{\frac{1}{2}}} e^{-\frac{1}{2}\left(\mathbf{x}-\boldsymbol{\mu}\right)^T\mathbf{\Sigma}^{-1}\left(\mathbf{x}-\boldsymbol{\mu}\right)} \tag{3.6}$$

Here, $\mathbf{x} \in \mathbb{R}^d$ denotes a point in a d dimensional Cartesian space, while $\boldsymbol{\mu} \in \mathbb{R}^d$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ correspond to the mean vector and covariance matrix of the Gaussian distribution. For our purposes, $d = 3$, $\boldsymbol{\mu}$ describes the mean position of the hand, and $\boldsymbol{\Sigma}$ captures covariance in hand position.

3.4 PDF's in Pose Space

Hand pose is represented as a joint probability distribution over position and orientation. We assume that the position and orientation of the hand are independent within a single cluster. Since there are two choices for the orientation component (Dimroth-Watson and girdle), the joint distribution takes on one of the following forms:

$$g\left(\mathbf{x}, \mathbf{q} | \theta\right) = p\left(\mathbf{x} | \theta_p\right) f\left(\mathbf{q} | \theta_f\right) \tag{3.7}$$

and

$$\bar{g}\left(\mathbf{x}, \mathbf{q} | \bar{\theta}\right) = p\left(\mathbf{x} | \theta_p\right) \bar{f}\left(\mathbf{q} | \theta_{\bar{f}}\right), \tag{3.8}$$

where θ and $\bar{\theta}$ are the parameters for the two joint distributions, θ_p consists of the parameters for the position density, and θ_f and $\theta_{\bar{f}}$ are the parameters of a Dimroth-Watson and girdle distributions, respectively.

The $g\left(\right)$ and $\bar{g}\left(\right)$ density functions essentially encode two different grasp types. The first constrains all six degrees of freedom of hand pose, while the second eliminates one of the orientation constraints.

3.5 Mixtures of Pose Models

An individual hand pose distribution can capture a single cluster of points, but a set of grasps is typically fit best by multiple clusters. Furthermore, the use of multiple clusters captures any covariance that may exist between the position and orientation of the hand when grasping a particular object. We therefore employ a mixture model-based approach. Here, the density function of the mixture, $h\left(\right)$, is defined as:

$$h(\mathbf{x}, \mathbf{q} | \boldsymbol{\Psi}) = \sum_{j=1}^{M} w_j c_j(\mathbf{x}, \mathbf{q} | \theta_j), \tag{3.9}$$

$$\boldsymbol{\Psi} = (w_1, ..., w_M, \theta_1, ..., \theta_M), \tag{3.10}$$

and

$$\sum_{j=1}^{M} w_j = 1, \tag{3.11}$$

where M denotes the number of component densities, and c_j is one of the two density functions describing hand pose ($g\left(\right)$ or $\bar{g}\left(\right)$). Each element of the mixture represents a single cluster of points, and is weighted by w_j. Estimation of the weights is accomplished using the Expectation Maximization algorithm (Dempster et al. 1977), and the parameters of the individual distributions are found by maximum likelihood estimation.

Let α_{ji} denote the probability that sample i belongs to cluster j, with the constraint that $\forall i \ \sum_{j=1}^{M} \alpha_{ji} = 1$. The first step in our EM process randomly assigns samples to clusters. For sample i, this is equivalent to setting α_{ji} to 1 for a randomly selected cluster whose index is j. Next, the parameters of each cluster are estimated using weighted versions of their maximum likelihood estimators (M-step). Note that in equation 3.9, θ_j denotes the tuple of parameters for cluster j, and corresponds to either $\left(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j, \mathbf{u}_j, k_j\right)$ or $\left(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j, \mathbf{u}_{1j}, \mathbf{u}_{2j}, k_j\right)$. The form of θ_j depends upon whether or not a Dimroth-Watson or girdle distribution is used to model orientation. The weighted maximum likelihood estimates of $\boldsymbol{\mu}_j$ and $\boldsymbol{\Sigma}_j$ are:

$$\boldsymbol{\mu}_j = \frac{\sum_{i=1}^{N} \alpha_{ji} \mathbf{x}_i}{\sum_{i=1}^{N} \alpha_{ji}}$$

and

$$\boldsymbol{\Sigma}_j = \frac{\sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{x}_i - \boldsymbol{\mu}_j\right) \left(\mathbf{x}_i - \boldsymbol{\mu}_j\right)^T}{\sum_{i=1}^{N} \alpha_{ji}}.$$

The axes of the Dimroth-Watson and girdle distributions are derived from the weighted sample covariance matrix:

$$\boldsymbol{\Lambda}_j = \frac{\sum_{i=1}^{N} \alpha_{ji} \, \mathbf{q}_i \mathbf{q}_i^T}{\sum_{i=1}^{N} \alpha_{ji}}.$$

Again, the MLE of $\mathbf{u}_j$ is parallel to the first eigenvector of $\boldsymbol{\Lambda}_j$, while the orthogonal vectors $\mathbf{u}_{1j}$ and $\mathbf{u}_{2j}$ span the same space as the first and second eigenvectors of $\boldsymbol{\Lambda}_j$. The maximum likelihood estimates of k_j for the Dimroth-Watson and girdle distributions uniquely satisfy:

$$G\left(k_j\right) \equiv \frac{F'\left(k_j\right)}{F\left(k_j\right)} = -\frac{\sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}\right)^2}{\sum_{i=1}^{N} \alpha_{ji}} \tag{3.12}$$

and

$$\bar{G}\left(k_j\right) \equiv \frac{\bar{F}'\left(k_j\right)}{\bar{F}\left(k_j\right)} = -\frac{\sum_{i=1}^{N} \alpha_{ji} \left[\left(\mathbf{q}_i^T \mathbf{u}_1\right)^2 + \left(\mathbf{q}_i^T \mathbf{u}_2\right)^2\right]}{\sum_{i=1}^{N} \alpha_{ji}}. \tag{3.13}$$

In the E-step, each α_{ji} and w_j is recomputed using the updated set of parameters as follows:

$$\alpha_{ji} \leftarrow \frac{w_j c_j\left(\mathbf{x}_i, \mathbf{q}_i | \theta_j\right)}{h\left(\mathbf{x}_i, \mathbf{q}_i | \mathbf{\Psi}\right)}$$

and

$$w_j \leftarrow \frac{\sum_{i=1}^{N} \alpha_{ji}}{\sum_{i=1}^{N} \sum_{j=1}^{M} \alpha_{ji}} = \frac{\sum_{i=1}^{N} \alpha_{ji}}{N}.$$

Finally, the expectation and maximization steps are repeated a finite number of times (in our case 30).

Chapter 4

Learning Grasp Affordances Through Human Demonstration

4.1 Problem Definition and Algorithm

In this chapter, we focus on the problem of describing the position and orientation of the hand as it approaches the object. Individual clusters are represented by a joint probability distribution over position and orientation. In our experiments, a single demonstration trial consists of a human teacher haptically exploring an object, pausing briefly in configurations that correspond to quality grasps. Experimental results demonstrate the feasibility of extracting a compact set of canonical grasps from the human demonstration. The learned representation can then be used to parameterize controllers that are capable of driving a hand to an appropriate pose for grasping, or to interpret the actions of other agents in the environment.

4.2 Data Collection Procedure

A human teacher wears a P5 glove (Essential Reality, Inc.) equipped with a Polhemus Patriot sensor near the wrist (Polhemus, Inc.).[1] These components continuously capture the pose of the hand at $15Hz$. A fixed transformation relative to the wrist is used to estimate the point between the tips of the thumb and index finger. Because the human teacher primarily uses precision grasps, this point is used as a description of hand pose. In addition, a Polhemus sensor is mounted on the object, which allows

[1] The experimental protocol was approved by the University of Oklahoma Internal Review Board (IRB #11909).

us to compute the pose of the hand in an object centered coordinate frame. Each trial consists of approximately 5 minutes of haptic exploration of the object. Throughout the course of a trial, the object may be translated and rotated in the global coordinate frame. This allows the human teacher to execute grasps that might not be possible if the object were in a fixed location of the workspace. During the trial, the teacher largely maintains contact with the object in configurations that correspond to quality grasps, although some samples fall along transitions between valid grasps. After the trial, the observations are subsampled and split into a training set and two validation sets. A training set consists of 1000 samples, each validation set contains 250 samples, and a total of 10 trials are performed for each object.

4.3 Model Selection

For a given set of observations, it is unclear *a priori* how many or of what type of cluster is appropriate. Our approach is to construct all possible mixtures that have a maximum of M clusters (we choose $M = 10$) and to choose the mixture that best matches the observations. For this purpose, we make use of the Integrated Completed Likelihood (ICL) criterion (Biernacki et al. 2000) to evaluate and order the different mixture models. Like the Bayesian Information Criterion, ICL prefers models that explain the training data, but punishes more complex models. In addition, ICL punishes models in which clusters overlap one-another. These features help to select models that describe a large number of grasps with a small number of clusters. ICL is computed as follows:

$$ICL = -2 \sum_{i=1}^{N} \sum_{j=1}^{M} \hat{\alpha}_{ji} \left[\log \left(w_j c_j \left(\mathbf{x}_i, \mathbf{q}_i | \theta_j \right) \right) \right] + \nu \log \left(N \right), \tag{4.1}$$

where:

$$\hat{\alpha}_{ji} = \begin{cases} 1 \text{ if } j = \operatorname{argmax}_l \alpha_{li} \\ 0 \text{ if } j \neq \operatorname{argmax}_l \alpha_{li} \end{cases},$$

and

$$l \in \{1 \cdots M\}.$$

Here, α_{li} denotes the probability that sample i belongs to cluster l, and ν denotes the total number of degrees of freedom in the mixture. Using this formulation, a lower ICL is more desirable.

The first term in equation 4.1 captures how well the mixture explains a provided set of samples. Note that if sample i is not assigned to cluster j (i.e., $\hat{\alpha}_{ji} = 0$), then it does not contribute to the likelihood term. Because each sample is assigned to a cluster, mixture models with well separated clusters result in lower ICL values, and are considered to be better solutions. The second term in the equation is a penalty term that takes into account the number of model parameters and the number of samples.

Because the EM algorithm is a gradient ascent method in a likelihood space containing many local maxima, each candidate mixture model was fit O times using the available training data (for our purposes, $O = 80$). For a given mixture, this ensures that a variety of different initializations for the EM algorithm are explored. The model that performs best on the first validation set according to ICL is subsequently evaluated and compared with other mixtures using the second validation set (again using ICL).

Due to our data collection procedure, some samples do not correspond to quality grasps, and instead correspond to transitions between grasps. It is desirable that our clustering algorithm be robust to this form of noise. However, when a large enough number of mixture components is allowed, the EM algorithm tends to allocate one or more clusters to this small number of "outlier" samples. We explicitly discard these mixture models when an individual cluster covers a very small percentage of the samples. In particular, a model is discarded when:

$$\frac{\max_j(w_j)}{\min_j(w_j)} \geq \lambda, \tag{4.2}$$

where λ is a threshold. For our experiments, we chose $\lambda = 5$ because it tends to result in the selection of high quality, compact models. Of the models that have not been removed by this filter step, the one with the lowest ICL measure on the second validation set is considered to be the best explanation of the observed data set.

4.4 Assessing Model Quality

In general, it is difficult to assess the performance of unsupervised learning techniques because there is no inherent notion of a ground truth. Thus, to assess the quality of a model that is produced by our clustering algorithm, we compare the learned clusters

to a set of heuristically chosen clusters. Our heuristic is based on knowledge of the object and the types of grasps demonstrated by the human teacher. For example, when grasping objects with handles, such as the heat gun shown in figure 4.1(d), a single grasp type is usually employed. The orientation of the hand relative to the object is typically fixed in a configuration orthogonal to the handle's major axis, and the hand position tends to cluster around the handle's center. This approach provides a reasonable estimate of the number of clusters that exist for an object, as well as an expectation for each cluster's shape or type.

Performance of the clustering algorithm is quantified in terms of a contingency table that counts the number of "true" positives (TP), "false" positives (FP), and "false" negatives (FN) present in a solution produced by our clustering algorithm. A true positive is scored when the algorithm identifies a cluster that corresponds to a heuristically derived cluster. A match occurs when the position component of a cluster covers the appropriate region of an object, and the orientation component captures the set of hand orientations used by the human teacher (i.e. whether or not a rotational symmetry exists). A false positive is scored when the algorithm identifies a cluster that does not match the heuristic. This can happen when the algorithm identifies multiple clusters where a single cluster should have been found. A false negative is scored when the algorithm fails to identify one of the heuristically chosen clusters.

Once a contingency table has been constructed for a grasp affordance model its true positive rate ($TPR = TP/(TP + FN)$), precision ($PRC = TP/(TP + FP)$), and false discovery rate ($FDR = FP/(TP + FP)$) are computed. The true positive rate reports the fraction of the desired clusters that were correctly identified, while the precision describes the fraction of correctly identified clusters out of the clusters actually learned by the algorithm. The false discovery rate is a measure of how much our algorithm is overfitting.

4.5 Experimental Results

In order to illustrate the capabilities of our clustering approach, we perform multiple grasping experiments using a variety of objects (see figure 4.1). Each object has its

(a) (b) (c) (d)

Figure 4.1: The set of objects used in the clustering experiments. (a) Cylinder; (b) Spray bottle; (c) Hammer; (d) Heat gun.

own unique set of grasps that may be modeled as a mixture of joint distributions over the position and orientation of the hand.

4.5.1 Cylinder

First, we consider the cylinder shown in figure 4.1(a). A total of four feasible grasps exist for this object: One for each end, and two along the major axis (corresponding to "overhand" and "underhand" configurations). Figures 4.2(a) and 4.2(b) depict a typical set of samples collected for the cylinder. Samples near the ends of the cylinder correspond to times when the hand is approximately aligned with and arbitrarily rotated about the object's major axis. Samples along the major axis occur when the hand is exploring the lateral surface of the cylinder in either an overhand or underhand configuration.

In figure 4.2(a), the 3D position of the hand is shown throughout the course of the experiment, while figure 4.2(b) provides a visualization of the corresponding hand orientations. Orientation of the hand is represented as a single point on the surface of the unit sphere: imagine that the object is located at the origin of the sphere; the point on the surface of the sphere corresponds to the intersection of the palm with the sphere. Note that this visualization technique aliases the set of rotations about the line perpendicular to the palm. For example, in figure 4.2(b), there is no way to distinguish grasps using an overhand configuration from those that use an underhand configuration.

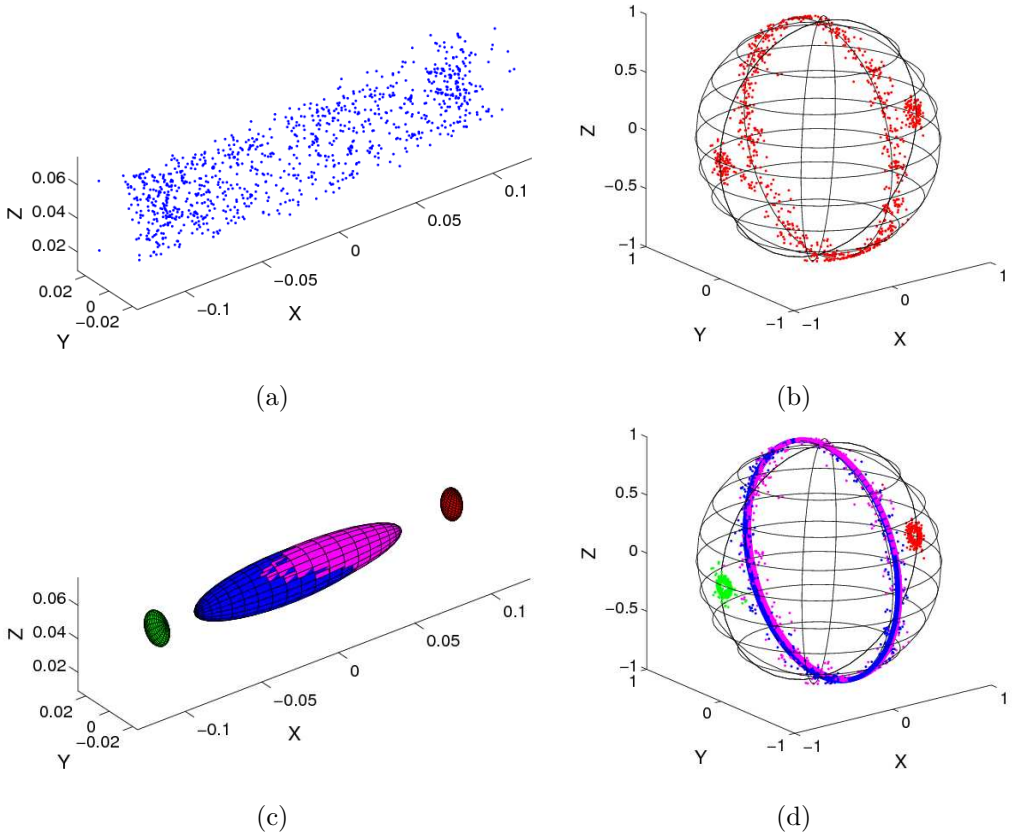

Figure 4.2: The training examples and learned affordance model for the cylinder. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

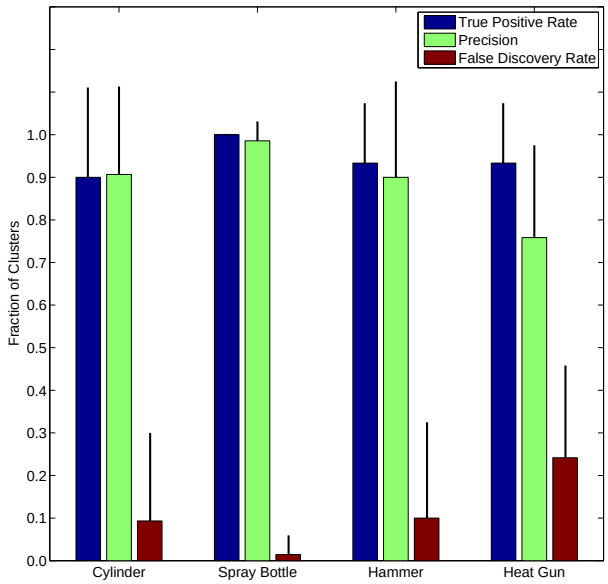

Figure 4.3: Contingency table summary for each of the objects used in the clustering experiments.

Figures 4.2(c) and 4.2(d) show the most common solution discovered by our algorithm for the cylinder. The position and orientation components of the model are shown independently, with similarly colored ellipsoids and circles constituting a single cluster in the pose space of the hand. Each circle represents the great circle on S3 defined by a girdle distribution, while each ellipsoid corresponds to the first standard deviation boundary of the corresponding multivariate Gaussian.

A single cluster for each grasp type was learned. Due to the significant rotational symmetries present in the object, a girdle distribution was selected as the orientation component for each grasp. The end grasps correspond to the green and red clusters, while grasps along the major axis are represented by the blue and magenta clusters. Because all positions along the major axis of the cylinder are viable for grasping, the corresponding ellipsoids have a larger volume, and are more elongated than those representing the end grasps. Also, notice that the position components of the overhand and underhand grasps overlap significantly. However, the algorithm selects two clusters to represent them because their corresponding orientation components are best

described by two different girdle distributions (even though our visual representation of orientation aliases this fact).

Figure 4.3 shows the mean true positive rate, precision, and false discovery rate for a variety of objects over the course of 10 trials. Focusing our attention on the cylinder, we see that on average the true positive rate and precision are above 0.9. Thus, a majority of the learned models matched our heuristic. However, there is a slight overfitting issue. As an example, on one occasion the algorithm learned three clusters for one of the grasps along the major axis of the cylinder, each with a Dimroth-Watson component for orientation, instead of a single cluster with a girdle distribution as its orientation component.

4.5.2 Spray Bottle

The next object we consider is the spray bottle shown in figure 4.4. Four of the feasible grasps are shown in the figure, with three of them having symmetric grasps achieved by rotating the object 180 degrees about its major axis. There is one grasp that may be made near the trigger (a), two from the side (b), two from the top (c), and two from the bottom (d). Each of these possibilities was extensively explored by the human teacher, and the collected training examples are shown in figures 4.5(a) and 4.5(b).

The affordance model learned by our algorithm for the spray bottle is shown in figures 4.5(c) and 4.5(d). Notice that there are a total of seven clusters (one for each demonstrated grasp type), where each of the colored line segments represent the mean rotation vector of a Dimroth-Watson distribution. The red and gray clusters represent the two symmetric grasps near the bottom of the spray bottle shown in figure 4.4(d). The green and brown clusters describe the set of grasps along the object's major axis as it is approached from the right and left, respectively (b). Notice that hand orientation is approximately orthogonal to the spray bottle's major axis in both cases. Also, the elongated nature of the ellipsoids captures the large variation in hand position along the major axis. Grasp approaches from the top of the spray bottle are represented by the orange and magenta clusters (c). Because the nozzle of the spray bottle is much smaller in comparison to its base, hand position and orientation are more constrained. This can be seen by comparing the relative volumes of the ellipsoids representing grasps near the base with those near the nozzle. Finally, the blue cluster

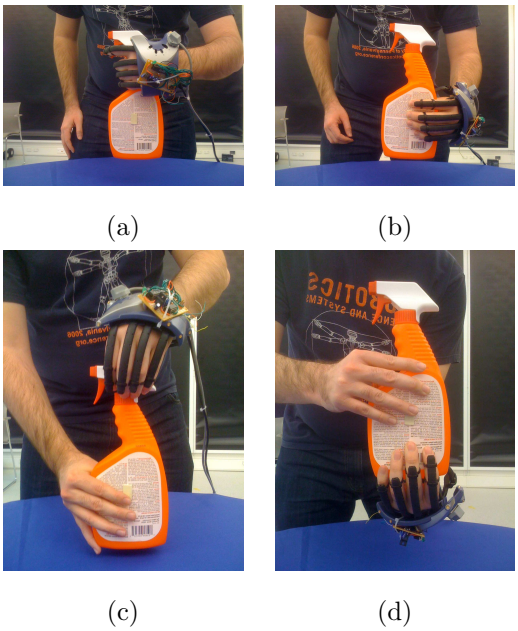

Figure 4.4: Four of the feasible grasps demonstrated by the human teacher for the spray bottle. (a) The trigger grasp; (b) The grasp from the right; (c) The top grasp; (d) The bottom grasp. Note that the grasps shown in (b), (c), and (d) have symmetric grasps that are achieved by rotating the spray bottle 180 degrees about its major axis.

captures the trigger grasp. Notice that in figure 4.5(a) the set of hand positions used to grasp the trigger seem to be comprised of two distinct sets of points. However, because the orientation of the hand does not vary much for grasps involving trigger, the algorithm allocates only a single cluster. This demonstrates the ability of our algorithm to generalize.

Over the ten experiments an average true positive rate of 1.0 was achieved. Therefore, every cluster we heuristically identified was in fact learned by the algorithm in each of the experiments. However, notice that the precision is slightly below 1.0. This is a result of overfitting: the algorithm split one cluster into two in a single experiment.

4.5.3 Hammer

The third object presented to the human teacher was the hammer shown in figure 4.1(c). This object was selected for a variety of reasons. First, it presented an

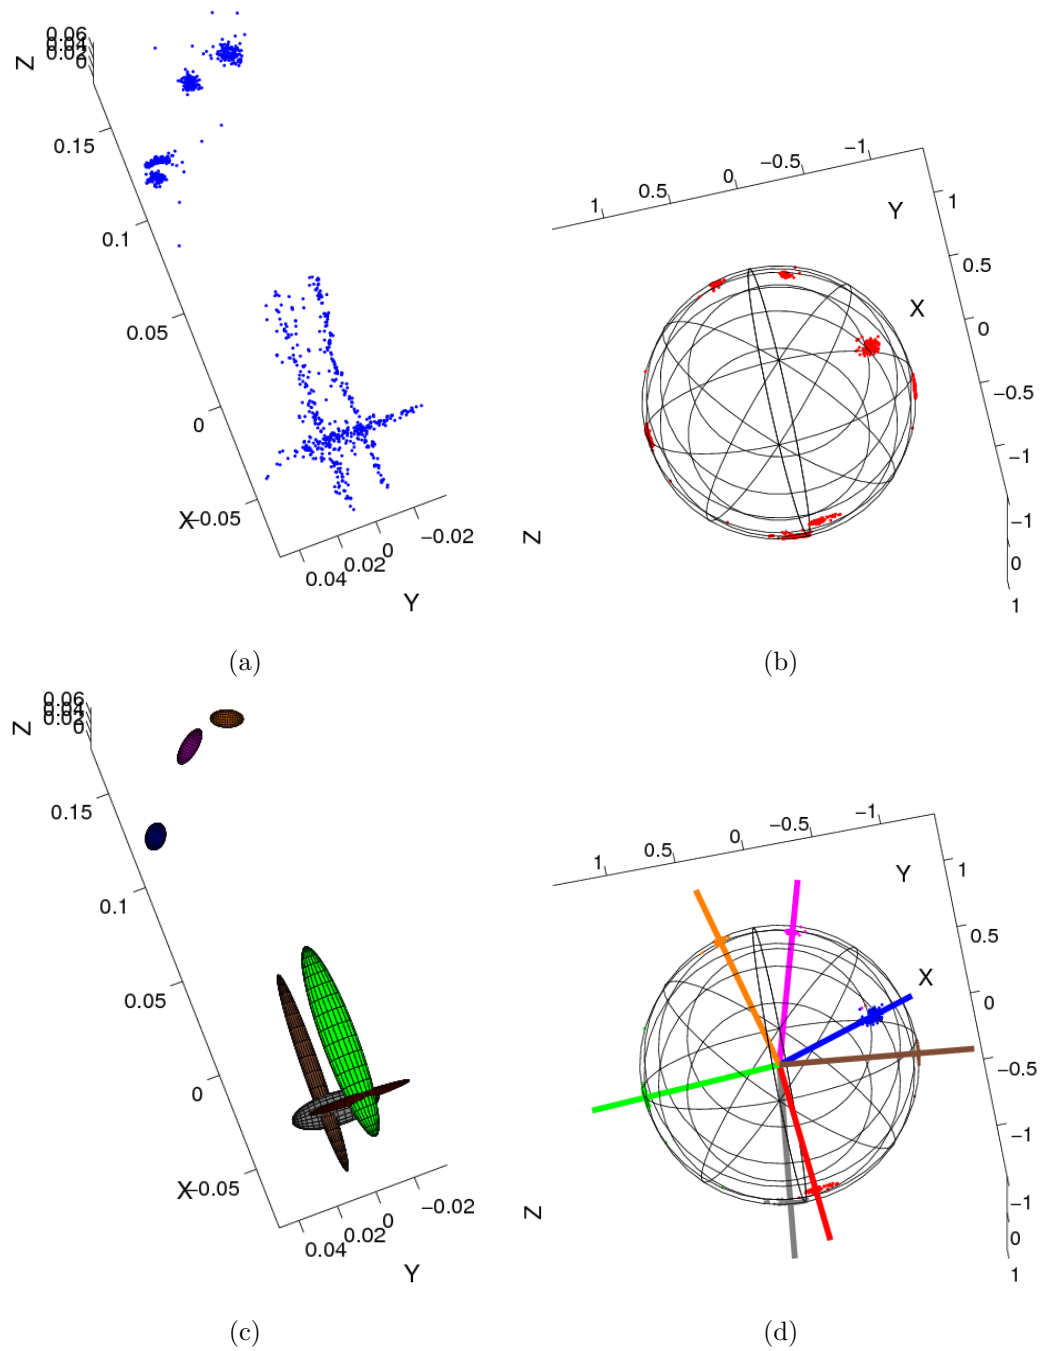

Figure 4.5: The training examples and learned affordance model for the spray bottle. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

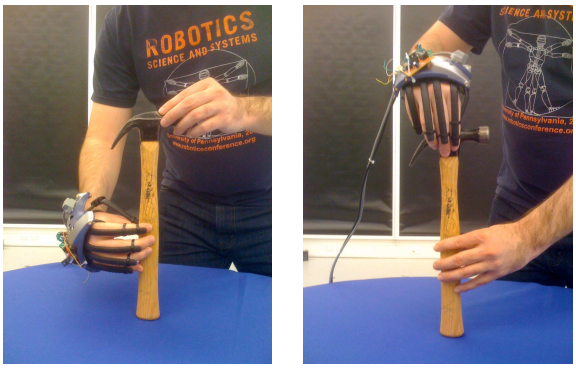

(a) (b)

Figure 4.6: The feasible grasps demonstrated by the human teacher for the hammer. (a) The handle grasp; (b) The head grasp.

interesting mixture of orientation constraints. Second, a hammer might be useful to a robot performing a real world task such as building a structure. The possible grasps include those near the handle and the head of the hammer, each of which was demonstrated by the human teacher, and are shown in figures 4.6(a) and 4.6(b), respectively. Note that the hand orientations used to grasp the handle always resulted in the thumb being closer to the head of hammer, and that the grasp shown in figure 4.6(b) has a symmetric grasp that is achieved by rotating the hammer 180 degrees about its major axis. A typical collection of training examples for these grasps is shown in figures 4.7(a) and 4.7(b).

The corresponding model learned by our algorithm is shown in figures 4.7(c) and 4.7(d). The green cluster represents the set of grasps that may be made with respect to the handle of the hammer. The elongation of the green ellipsoid along the major axis essentially encodes the handle's length. This means that any position on it is viable for grasping. In addition, the hammer may be grasped as long as the orientation of the hand is approximately orthogonal to the handle. Thus, a girdle distribution was learned for this portion of the object. The other two clusters (red and blue) involve grasps near the head of the hammer. These grasps are symmetric in that they are accomplished by grasping the object from the top, rotating the hammer 180 degrees, and then re-grasping. Since the head of the hammer is much smaller in comparison to its handle, the red and blue ellipsoids have less volume than the green ellipsoid. Also, because grasping this portion of the object involves very constrained

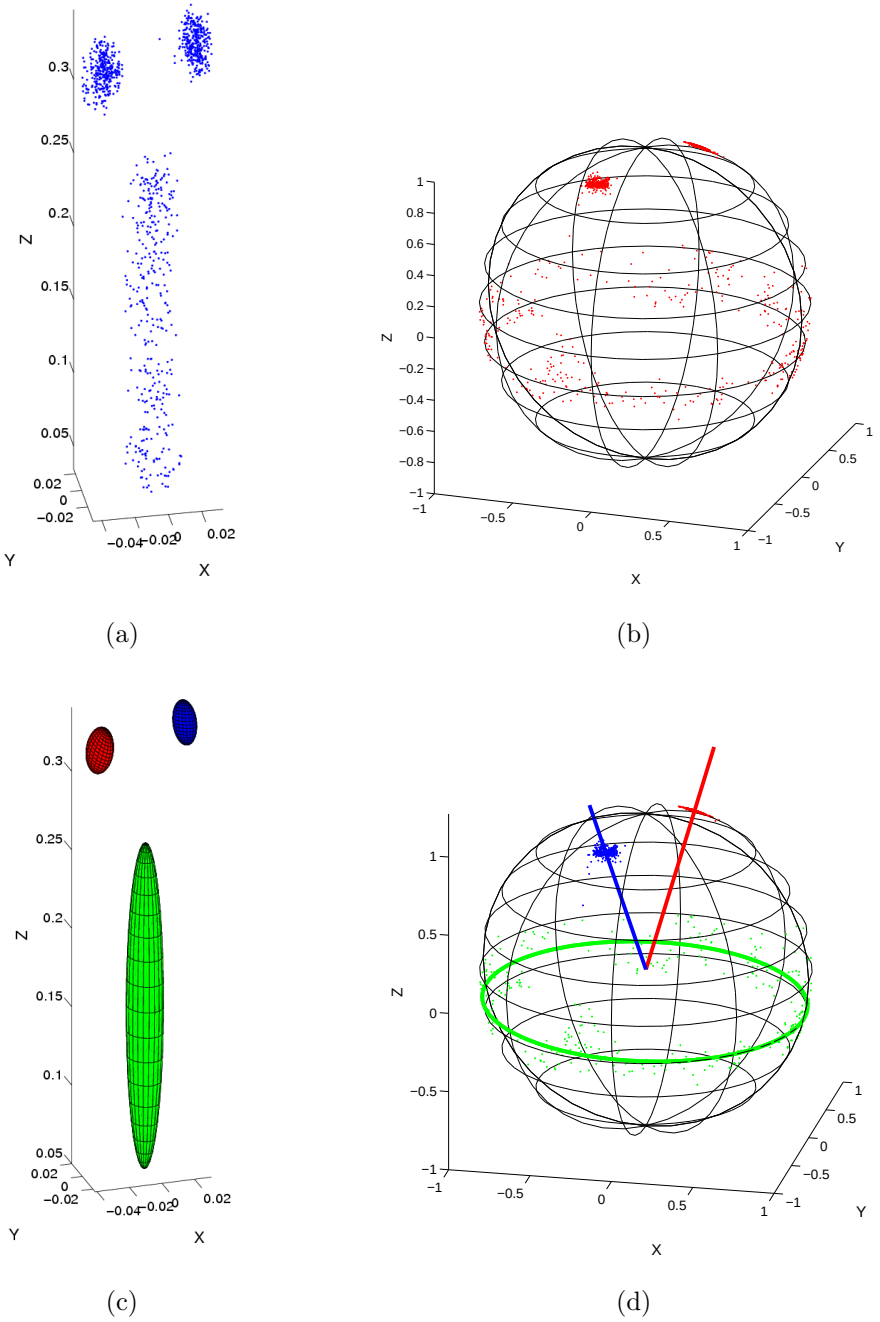

Figure 4.7: The training examples and learned affordance model for the hammer. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

(a) (b) (c)

Figure 4.8: The feasible grasps demonstrated by the human teacher for the heat gun. (a) The handle grasp; (b) The nozzle grasp; (c) The top grasp.

hand orientations, the algorithm selected Dimroth-Watson distributions to model them.

4.5.4 Heat Gun

The heat gun shown in figure 4.1(d) was selected for many of the same reasons as the hammer. It has a variety of orientation constraints, and is similar in shape to other real world objects such as a drill. The feasible grasps for this object are shown in figure 4.8. In these experiments, the portion of the heat gun that is on the other side of the handle was not grasped by the human teacher. Furthermore, the hand orientations used to grasp the heat gun's nozzle always resulted in the thumb being closer to the handle (see figure 4.8(b)).

Even though three primary grasps were demonstrated, our algorithm tends to choose solutions with four clusters (figures 4.9(c) and 4.9(d)). This is reflected by the higher false discovery rate shown in figure 4.3. The red cluster corresponds to grasps along the handle of the heat gun. A Dimroth-Watson distribution is used to model the orientation of the hand because the human teacher only grasped the handle in such a way that affords the use of the trigger. The blue cluster represents approaches from the top where the hand may be arbitrarily rotated about the object's major axis, with position being relatively constrained. The remaining clusters (green and magenta) describe the set of grasps along the lateral surface of the heat gun's nozzle. Ideally this would be a single cluster, but the algorithm preferred to separate it into two. This may be due to the fact that the nozzle widens as one approaches the handle, which means the distribution in position varies more in this region. However, both clusters did capture the rotational symmetry in hand orientation afforded by

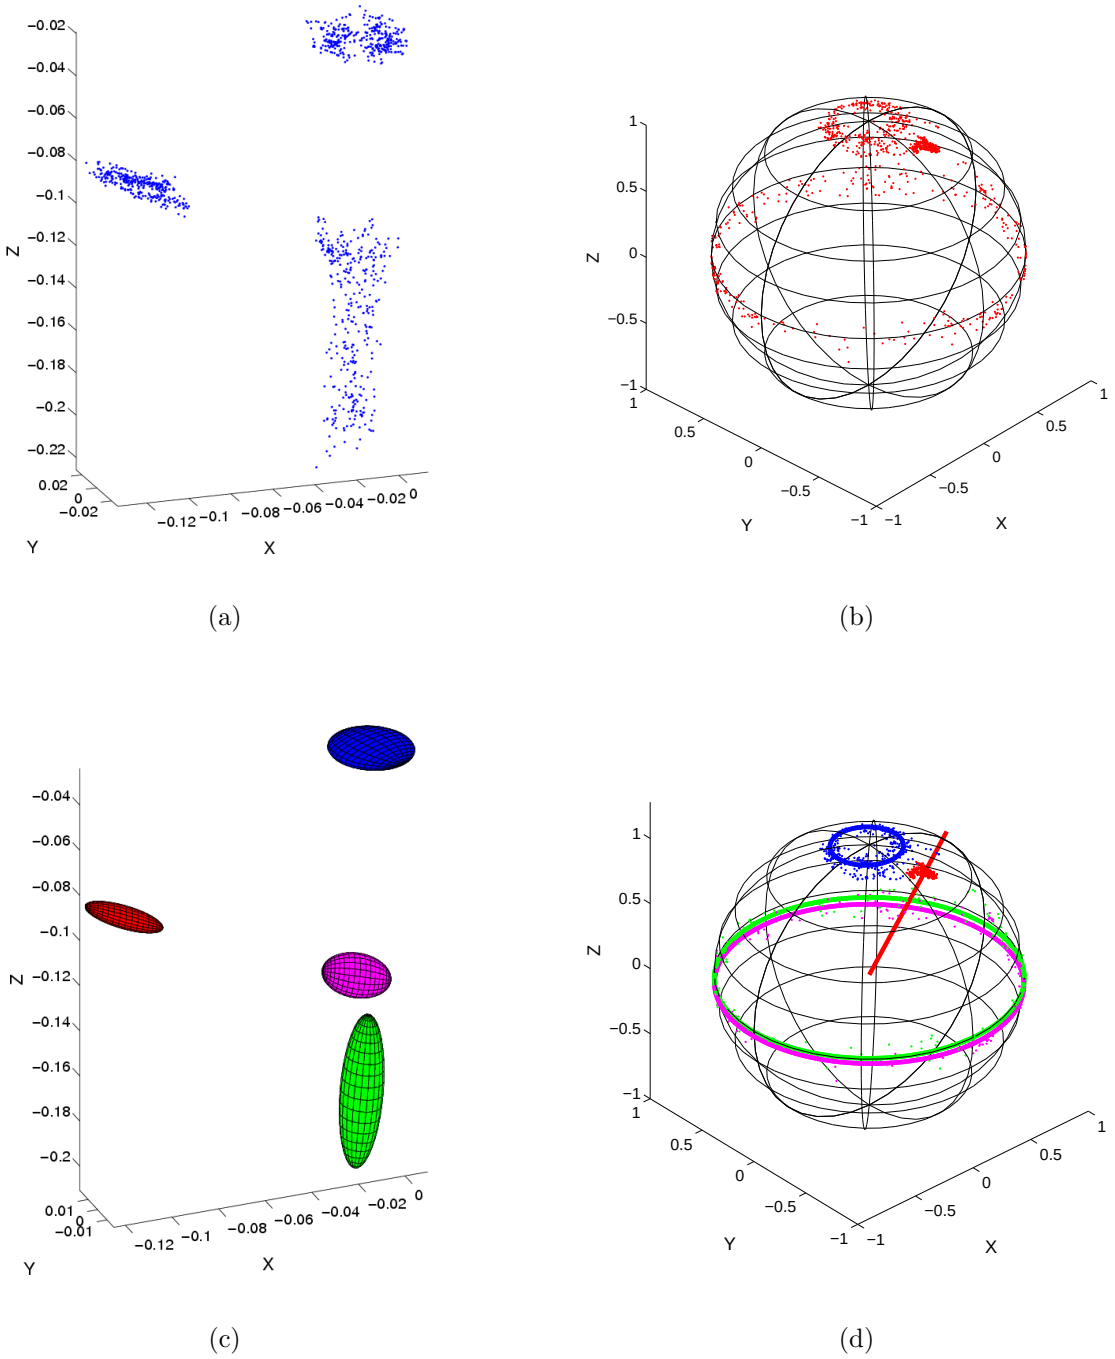

Figure 4.9: The training examples and learned affordance model for the heat gun. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

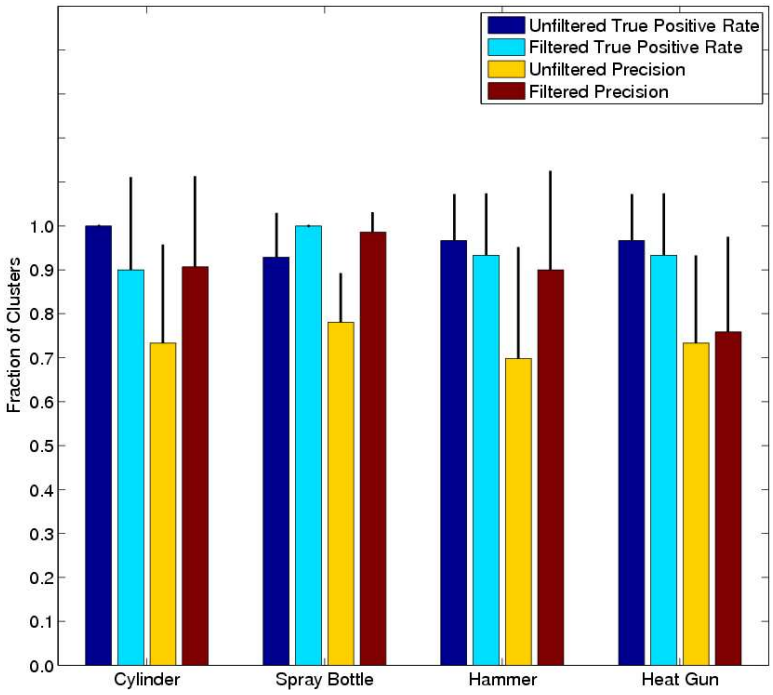

Figure 4.10: Contingency table comparison of the unfiltered and filtered versions of the grasp affordance learning algorithm for each of the objects used in the clustering experiments.

the nozzle. Thus, while not the expected solution, the algorithm learned a reasonable solution a majority of the time. This is supported by the high true positive rate exhibited by the heat gun.

This result for the nozzle may be due to several different factors. First, Gaussian distributions may not be suitable for modeling the positions of the hand when grasping objects exhibiting a conical geometry. Second, the algorithm may also be overfitting. Finally, it is possible that the heuristic is incorrect, and the algorithm is actually finding an underlying structure in the pose space of the hand.

4.6 Sensitivity Analysis

In order to gain a better understanding of the grasp affordance learning algorithm, its performance is compared to that of a baseline algorithm which removes the filtering step described in section 4.3. Figure 4.10 shows the mean true positive rate and mean

precision for the filtered and unfiltered versions of our algorithm. Notice that the mean true positive rate drops slightly for most of the objects when filtering is introduced. This is typically due to the elimination of solutions where a learned cluster matches one of the heuristically defined clusters, but it explains a small number of training examples.

However, for the spray bottle, the absence of filtering causes more false negatives to be scored because girdle distributions are learned as the orientation component of a cluster, even though Dimroth-Watson distributions were specified by the heuristic. When filtering is introduced, the algorithm tends to throw out these solutions, and as a result, the true positive rate increases.

Notice that for each of the objects the mean precision is lower when no filtering is performed. This occurs because the unfiltered version of the algorithm tends to allocate more clusters than the filtered version of the algorithm. While many of these clusters may match the heuristic (explaining the high mean true positive rates), some of them are unnecessary. When filtering is performed, many of the solutions with extra clusters are eliminated, which results in a higher mean precision for each of the objects.

When all of the objects are considered together, the mean decrease of the true positive rate from the unfiltered to the filtered version of the algorithm is 0.02 ($p < 0.36$ according to a paired bootstrap test). However, the overall effect of introducing filtering produces a mean increase in precision of 0.15 ($p < 10^{-4}$). This suggests that filtering facilitates the selection of models that better match our heuristic.

4.7 Discussion

In this chapter, we have presented a technique for learning canonical hand positions and orientations for reach-to-grasp actions. Compact representations are constructed from many example grasps made by clustering the pose of the hand. For a given object, we want the set of affordances to be small. This property enables the use of affordances as a way to access "primitives" in higher-level activities, including planning, learning, and the recognition of motor actions by other agents (Brock et al. 2005; Fagg et al. 2004).

In particular, the clusters that have been learned map directly onto resolved-rate controllers that can bring a robot hand to a specific position and orientation relative to the object. Note that this control step makes two assumptions: first, that the robot has a similar hand morphology to the human demonstrator; second, that haptic exploration methods are available to refine the grasps once the hand is approximately in the right configuration (Coelho and Grupen 1997; Platt et al. 2002). .

Chapter 5

Learning Grasp Affordances Through Robot Teleoperation

5.1 Problem Definition and Algorithm

The experience needed to learn the grasp affordances of an object does not have to be derived from the observation of human behavior. Rather, it may come from a robot automatically performing a grasping task, or under the guidance of a human teleoperator. In this chapter, a method that learns a mapping from objects to grasps through robot teleoperation is presented. Experiments are performed using NASA's humanoid robot *Robonaut*. A human teleoperator generates experience by haptically exploring a given object, and a small number of canonical grasps are associated with the object using the methods described in chapters 3 and 4.

5.2 Data Collection Procedure

The human teleoperator is able to control Robonaut's many degrees of freedom with a virtual reality-like helmet and a data glove equipped with a Polhemus sensor (Ambrose et al. 2000). In addition to articulating Robonaut's neck, the helmet provides visual feedback from the environment to the teleoperator. The arms and hands of the robot are commanded by tracking the movements of the human's wrists and fingers, and performing a mapping from human motion to robot motion. The origin of the coordinate frame that describes hand pose is located on the back of Robonaut's hand, with the positive x axis pointing in the direction of the finger tips when they are fully extended, and the positive z axis being orthogonal to the back of the hand.

Each trial consists of the human teacher haptically exploring an object for approximately 15 minutes. The object is located in a fixed pose relative to the robot. Hence, the object cannot be moved throughout the course of a trial. To maximize the number of quality samples collected, different grasping strategies may be employed by the teleoperator based on the local geometry of the object. For example, when grasping larger surfaces, a sliding motion in conjunction with a fixed finger configuration is used. This ensures that the feasible positions and orientations of the hand are collected in a timely manner. In contrast, the teleoperator repeatedly opens and closes the robot's hand when grasping small surfaces. This strategy forces hand pose to vary even though the hand may not be able to slide along the local surface.

When compared with the data collected during direct observation of a human performing grasping actions, the robot teleoperation experience tends to contain larger amounts of noise. Robonaut's arm motions are slower and less fluid under human control. Hence, more samples correspond to transitions between grasps. To alleviate this problem the transitions are removed manually by identifying the time intervals in which they occur. Figure 5.1 illustrates the result of this process. The blue samples correspond to times when Robonaut's hand is in contact with the object, whereas samples in red are labeled as transitions between grasps. The blue examples are then subsampled, and partitioned into a training set and two validation sets.

5.3 Experimental Results

In order to illustrate the capabilities of our clustering approach when applied to experience gathered using an actual robot, we perform multiple grasping experiments using a metallic bar and a hammer. Because our time with Robonaut was limited, we were only able to collect two trials of grasping experience for each object.

5.3.1 Bar

The first object we consider is the bar shown in figure 5.2(a). Notice that one of its ends is fixed to the table. As a result, there are only three feasible grasp types that may be made with respect to this object: one from above, one from below, and one from the side. Each of these possibilities was extensively explored by the

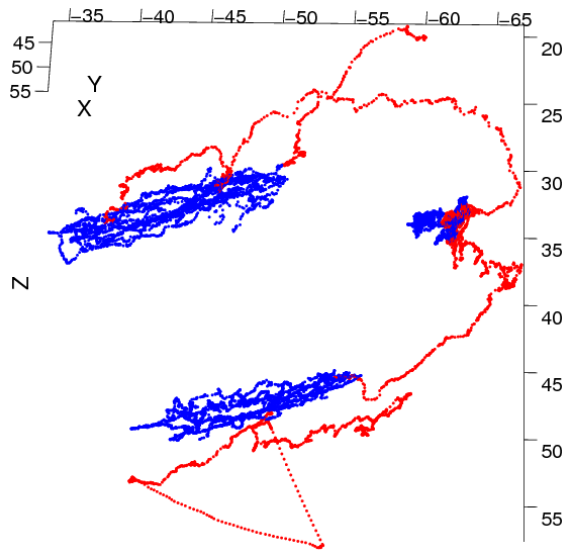

Figure 5.1: The position of Robonaut's hand throughout the course of a single experimental trial. The blue samples correspond to grasps, whereas the red samples correspond to transitions between grasps.

(a) (b)

Figure 5.2: The set of objects used in the Robonaut clustering experiments. (a) Bar; (b) Hammer.

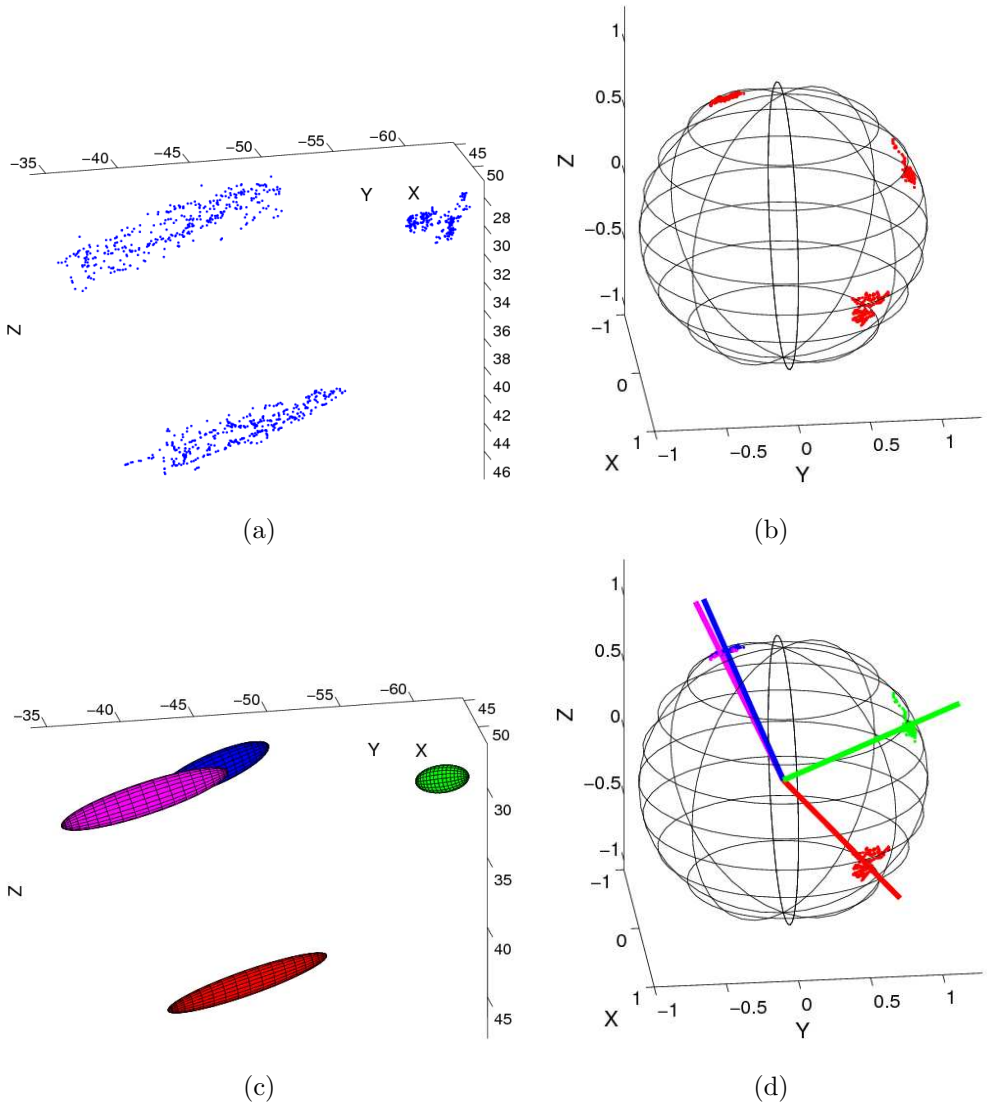

Figure 5.3: The training examples and learned affordance model for the bar. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

human teleoperator. Figures 5.3(a) and 5.3(b) show the examples that remained after eliminating all transitions between grasps.

Figures 5.3(c) and 5.3(d) show the most common affordance model learned by our algorithm. A total of four distinct hand pose clusters were discovered. The magenta and blue clusters correspond to the overhand grasp, the red cluster describes the underhand grasp, and the green cluster represents the side grasp.

First, notice that the magenta, blue, and red ellipsoids, which capture the viable positions of the hand when grasping the bar from above and below, are elongated in a direction that is parallel to the major axis of the bar. Also, the amount of variation along this direction is indicative of the bar's length. While this is an intuitive result, the algorithm splits the overhand grasps into two clusters. This may have been due to overfitting, but upon a closer examination of the training data, some of the samples near the end of the bar appear to shift towards the robot. Hence, the algorithm may have allocated an additional cluster to account for this change in hand position. The green ellipsoid represents the set of viable hand positions when grasping the bar from the side. Its volume is smaller than those of the magenta, blue, and red ellipsoids because hand position is more constrained for the side approach.

Because the surface of the palm tends to be parallel with the sides of the bar when it is grasped, the orientation components of the learned affordance model are represented by Dimroth-Watson distributions. Notice that the magenta and red vectors are approximately opposite in direction; this captures the relationship between the hand orientations used to grasp the bar from above and below. The green vector corresponds to the mean orientation used when approaching the bar from the side. Notice that is a 90 degree rotation away from the orientations used for the overhand and underhand grasp configurations.

5.3.2 Hammer

The next object we consider is the hammer shown in figure 5.2(b). There are two possible grasp approaches for this object: one from the top when grasping the head of the hammer, and one from the side when grasping its handle. The training examples collected for a single experiment are shown in figures 5.4(a) and 5.4(b).

While the learned affordance model shown in figures 5.4(c) and 5.4(d) is compact, it did not match our heuristic. Because the hand may be arbitrarily rotated

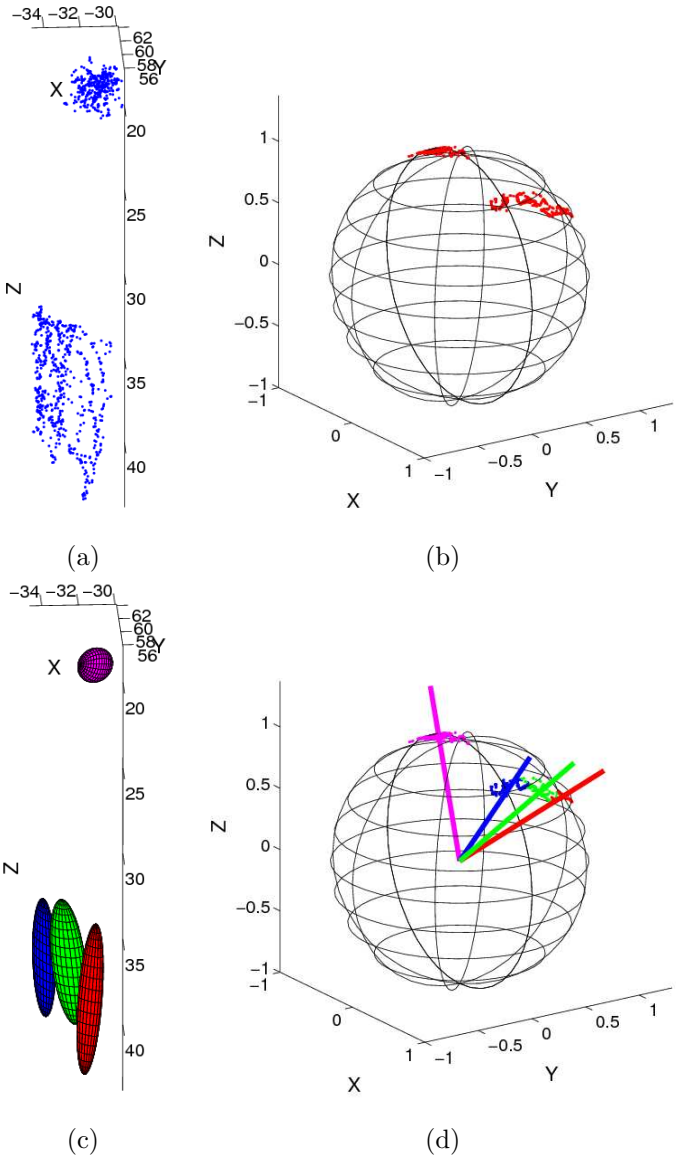

Figure 5.4: The training examples and learned affordance model for the hammer. (a) The position of the hand; (b) The orientation of the hand; (c) The position component of the learned affordance model; (d) The orientation component of the learned affordance model.

about the hammer's handle, one might expect the orientations used to grasp it would be represented best by a girdle distribution. However, the algorithm selected three Dimroth-Watson distributions instead. There are a couple of possible explanations for this. First, it was difficult for the teleoperator to rotate the hand of the robot all the way around the hammer's handle. Hence, there may not have been enough evidence of a rotational symmetry to justify the selection of a girdle distribution. Second, because a point on the *back* of the hand is used to measure where grasp events occur, there may have been sufficient separation in hand position to justify the use of three clusters. This claim is supported by the non-overlapping red, green, and blue ellipsoids.

The magenta cluster represents the canonical hand pose used when grasping the head of the hammer. Notice that the mean rotation of the cluster, which is represented by the magenta vector, is approximately orthogonal to the red, green, and blue vectors as one would expect. Also, the magenta ellipsoid has a much smaller volume than those ellipsoids representing the side grasp because hand position is more limited when grasping the hammer's head.

5.4 Discussion

In this chapter, a method that learns a mapping from objects to grasps through robot teleoperation was presented. A human teacher demonstrates the feasible grasps that may be made with respect to a particular object, and a small number of canonical grasps are associated with it by clustering the poses of the hand. Experimental results demonstrated the ability of our algorithm to learn compact representations for two different objects.

One question is: are the learned representations sufficient to parameterize grasping actions? In the case of the bar, the algorithm learned similar models for each experiment. However, both solutions did not match our heuristic: the viable overhand grasps were split into two clusters. While this was not ideal, the learned model does compactly describe the many different options available when grasping the bar. Furthermore, the learned representation should enable the robot to successfully grasp the object.

In the case of the hammer, the algorithm also learned similar models for each experiment, but the representations learned for each side approach were unexpected. Possible solutions to this problem include performing an additional transform from the back of the hand to a point that is expected to be closer to the object when grasping, and to track the pose of the object visually. The latter would enable the teleoperator to better explore the object because it could be translated and rotated throughout the workspace. More importantly, taking this step would be an initial attempt at addressing the perception problem one faces when taking an indirect approach to learning grasp affordances.

Another direction for future work is the planning and executing of two-handed grasps, sequences of re-grasps, or operations in which an object is handed from one agent to another. There are a few different approaches one could take in solving this problem. In the case of Robonaut, one could apply an affine transformation on the affordance model learned for the left arm to generate a model for the right arm (Sweeney and Grupen 2007). Alternatively, a more direct approach would be to define joint probability density functions over the pose of both hands, and perform clustering in the resulting two handed pose space.

Chapter 6

Grasp Affordances for Finger Configurations

6.1 Problem Definition and Algorithm

The focus of this work thus far has been on how an agent positions and orients its hand relative to an object when grasping. However, the configuration of the fingers has a great impact on whether or not a grasp will be successful. In this chapter, I present a method that learns the canonical hand preshapes used to grasp a given object. This is accomplished through the construction of joint probability density functions over hand position, hand orientation, and finger configuration. In order to address the high dimensionality of the finger configuration space, clusters are described in a low dimensional *eigengrasp* space. Experiments are performed using NASA's humanoid robot *Robonaut*. First, grasping experience for a number of objects is provided by a human teleoperator, and is used to learn the set of eigengrasps. Then, the teleoperator proceeds to grasp another set of objects, for which low dimensional representations of finger configuration are constructed. Finally, for each object a small set of canonical preshapes that model the ways in which the object may be grasped are learned.

6.2 Eigengrasps

Humanoid robots such as Robonaut typically have many degrees of freedom (DOF) available to perform manipulation tasks. For example, each of Robonaut's hands has 12 DOF: three for the thumb, index, and middle fingers; one for the ring and pinkie fingers; and one for the palm, which can help to grasp objects by cupping them (Ambrose et al. 2000). Incorporating finger configurations into our clustering algorithm

is a key step to constructing more complete grasp affordance representations. One possible approach to this problem is to learn clusters using the full dimensionality of the robot's end-effector. However, hands with a large number of joints can be difficult to model because as the dimensionality of a space increases, the Euclidean distance metric can become meaningless (Beyer et al. 1999), making it more difficult to find distinct clusters. Furthermore, an increasingly large number of training examples is needed to adequately sample a space as more and more dimensions are added to it.

One question is whether or not all of the available DOF of the hand are even necessary to accurately model the finger configurations used to grasp. For example, when executing a power grasp, the fingers tend to flex in unison. This means that there is a strong correlation between the distal and proximal joints of each finger, as well as a correlation across fingers. Santello et al. (1998) and Ciocarlie et al. (2007) present approaches that take advantage of such correlations through the notion of an *eigengrasp*. In the robotics domain, the eigengrasps of a manipulator are a set of basis vectors in the joint space of the hand. Linear combinations of a small number these eigengrasps can be used to approximate the finger configurations used when grasping.

More formally, let $\mathbf{p} \in \mathbb{R}^d$ be a column vector of joint angles describing the finger configuration of a robot's end-effector, and $\mathbf{V} \in \mathbb{R}^{d \times d}$ constitute a basis for the vector space of which $\mathbf{p}$ is a member. The columns of $\mathbf{V}$ represent directions in the joint space of the hand (the eigengrasps), and are ordered from those that capture the most variance in finger configuration to the smallest (i.e., from the largest corresponding eigenvalue to the smallest). Linear combinations of the columns of $\mathbf{V}$ can be used to represent any possible pose for the fingers of the robot's hand:

$$\mathbf{p} = \sum_{i=1}^{d} a_i \mathbf{v}_i = \mathbf{V}\mathbf{a}. \tag{6.1}$$

Here, $\mathbf{v}_i \in \mathbb{R}^d$ is the i'th column of $\mathbf{V}$, $\mathbf{a} \in \mathbb{R}^d$ is a column vector of coefficients, and $a_i \in \mathbb{R}$ is element i of the vector $\mathbf{a}$.

Because there may be a large number of joints in the robot's hand, the configuration of the fingers may be approximated by using a small number (K) of eigengrasps:

$$\mathbf{p} = \sum_{i=1}^{K} a_i \mathbf{v}_i = \hat{\mathbf{V}}\hat{\mathbf{a}}, \tag{6.2}$$

where:

$$\hat{\mathbf{V}} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \ldots & \mathbf{v}_K \end{bmatrix}, \tag{6.3}$$

and

$$\hat{\mathbf{a}} = \begin{bmatrix} a_1 & a_2 & \ldots & a_K \end{bmatrix}^T. \tag{6.4}$$

Given a finger configuration $\mathbf{p}$ and a subset of the eigengrasps $\hat{\mathbf{V}}$, a low dimensional representation of $\mathbf{p}$ is obtained by solving the system of linear equations in 6.2 for $\hat{\mathbf{a}}$.

One question yet to be answered is: how are the eigengrasps computed? Our approach is to learn them by performing principal component analysis on the finger configurations used by a teleoperator to grasp a variety of objects (Hand et al. 2001). Let $\mathbf{P} \in \mathbb{R}^{d \times N}$ be the set of finger configurations resulting from the human demonstration, where N denotes the number of examples. The eigengrasps are determined by computing the eigenvectors of $\mathbf{P}$'s covariance matrix.

6.3 Modeling Grasp Preshapes

To incorporate the finger configuration of the hand into the grasp affordance learning algorithm, joint distributions over hand pose and the low dimensional representation of finger configuration are constructed. Recall that the PDF's $g\,()$ and $\bar{g}\,()$ both represent variations around a canonical pose for the hand when grasping. Extending the approach to model grasp preshapes is accomplished by redefining $g\,()$ and $\bar{g}\,()$ as follows:

$$g\left(\mathbf{x}, \mathbf{q}, \hat{\mathbf{a}} | \theta\right) = p\left(\mathbf{x} | \theta_p\right) f\left(\mathbf{q} | \theta_f\right) p\left(\hat{\mathbf{a}} | \theta_h\right), \tag{6.5}$$

and

$$\bar{g}\left(\mathbf{x}, \mathbf{q}, \hat{\mathbf{a}} | \bar{\theta}\right) = p\left(\mathbf{x} | \theta_p\right) \bar{f}\left(\mathbf{q} | \theta_{\bar{f}}\right) p\left(\hat{\mathbf{a}} | \theta_h\right). \tag{6.6}$$

Here, $p\left(\hat{\mathbf{a}} | \theta_h\right)$ is a multivariate Gaussian distribution over K dimensions, and it is assumed that hand position, hand orientation, and finger configuration are conditionally independent given a cluster. In terms of hand pose, a cluster roughly corresponds to some portion of the object. Presumably, the hand poses used to grasp are similar due to smoothness of the object's local geometry. Hence, the portion of the object that corresponds to the identified cluster affords some class of finger configurations, but this class should not (as a first approximation) be dependent upon the position and orientation of the hand within the cluster.

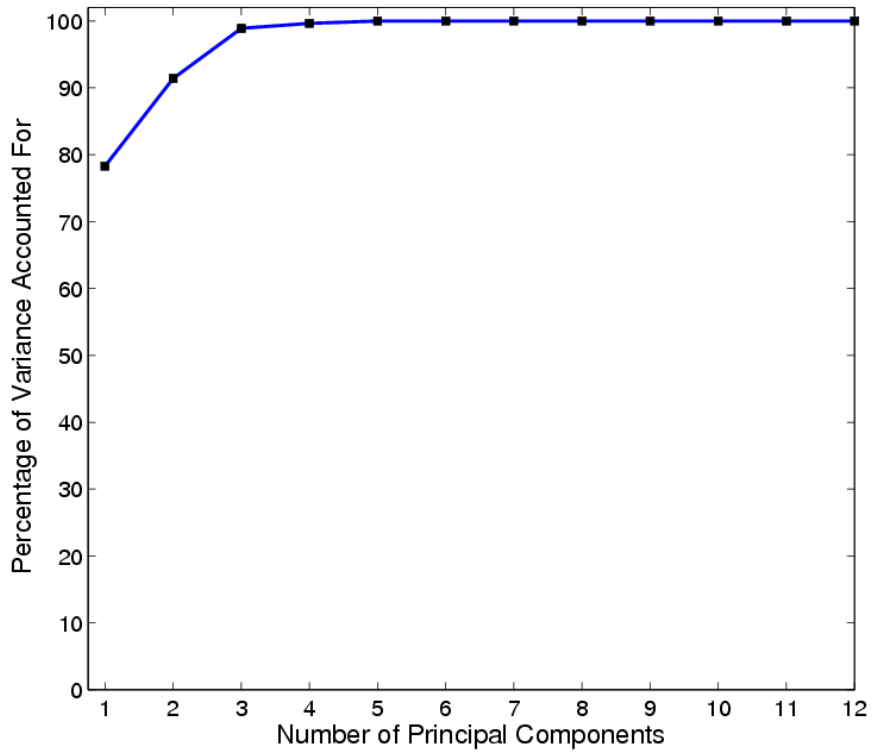

Figure 6.1: The percentage of the finger configuration variance accounted for by each principal component (eigengrasp).

6.4 Experimental Results

To demonstrate the effects of incorporating finger configuration into the grasp learning process, a number of experiments are performed. First, the eigengrasps are learned based on experience that is generated by a human teleoperator. A number of different objects are used to ensure a reasonable sampling of the finger configurations used to grasp. Then, low dimensional representations of the hand shapes used to grasp the bar and hammer presented in section 5.3 are constructed, for which a small number of canonical preshapes are learned.

6.4.1 Eigengrasps

Robonaut's eigengrasps are learned by concatenating the teleoperation data collected for a number of objects, and performing principal component analysis on the resulting set of finger configurations. Due to invalid sensor data, seven of the finger joints are

ignored. This means that the number of effective degrees of freedom in Robonaut's hand has been reduced from twelve to five.

Figure 6.1 shows the cumulative percentage of the finger configuration variance that is accounted for as the number of principal components is increased. Notice that the cumulative PVAF quickly approaches 100%. This suggests that there are many redundant degrees of freedom when grasping the set of objects used in the experiment. Furthermore, approximately 98% of the variance can be explained by the first three principal components alone. This, in conjunction with the ability to visualize the resulting low dimensional representation of finger configuration, led to the use of only the first three eigengrasps (i.e., $K = 3$).

6.4.2 Bar

Figure 6.2 shows the training examples and the learned affordance model for the bar. Notice that a total of five clusters were learned: two for the overhand configuration, two for the underhand configuration, and one for the side approach. In terms of hand pose, this result is very similar to the one presented in section 5.3.1.

Figure 6.2(f) shows the learned eigengrasp clusters. Each corner of the bounding box provides a visualization of the mapping that occurs between the low dimensional representation of finger configuration and each joint of the robot's hand. Notice that variation along the first eigengrasp corresponds to flexion of the index and middle fingers, while variation along the second eigengrasp causes adduction and abduction of the thumb. However, variation along the third eigengrasp does not seem to affect the configuration of the fingers significantly. Also, note that the ring and pinkie fingers remain in their extended configurations because they are among the degrees of freedom that have been zeroed out.

Turning to the learned eigengrasp clusters, notice that the green and red ellipsoids are in the same region of the finger configuration space even though they correspond to grasp approaches from above and below the bar. Because the same sliding technique was employed by the teleoperator when demonstrating these grasps, the hand had a similar shape for each approach. However, for the red cluster there is more variation in finger configuration, which explains the elongation of the red ellipsoid. In contrast, the hand was continually opened and closed when the side approach was used to grasp the bar. This is evident by comparing the hand shapes that correspond to

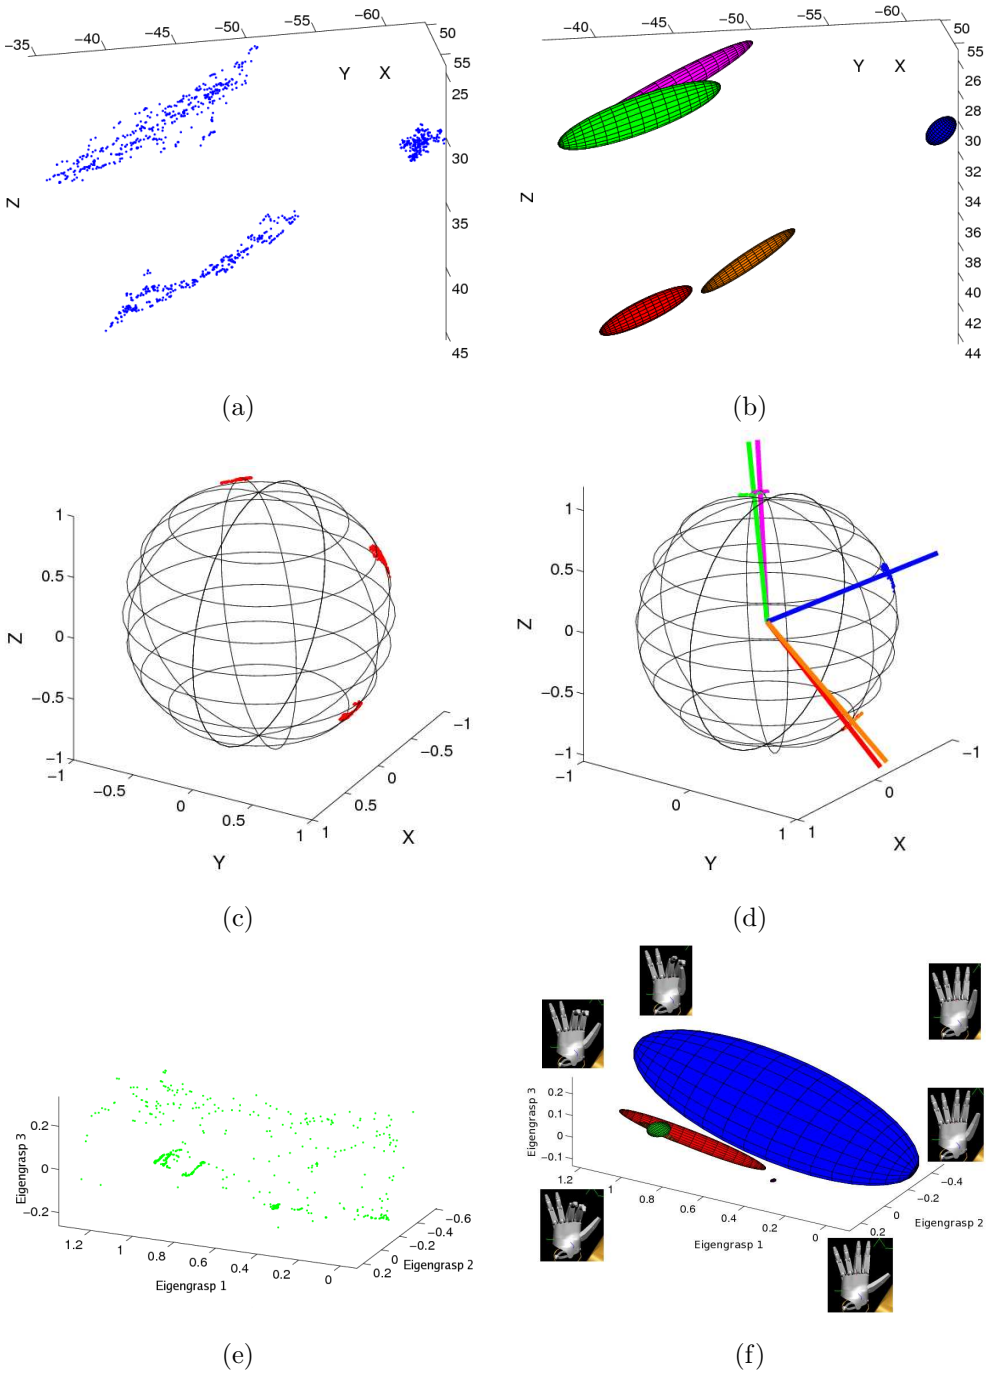

Figure 6.2: The training examples and learned affordance model for the bar. (a) The position of the hand; (b) The position component of the learned affordance model; (c) The orientation of the hand; (d) The orientation component of the learned affordance model; (e) The finger configuration of the hand; (f) The finger configuration component of the learned affordance model.

points on opposite ends of the blue ellipsoid's major axis. On the right end, the hand is in an open configuration, but on the left end the middle and index fingers are flexed considerably. Also, notice that the blue ellipsoid is separated from the other eigengrasp clusters, which highlights the different hand shapes used when grasping the bar from above and below versus from the side.

6.4.3 Hammer

The example grasps demonstrated by the human teleoperator and the learned grasp affordance model for the hammer are shown in figure 6.3. In this case five clusters were learned: the red cluster represents grasps when approaching from above the hammer's head, while the remaining clusters capture grasps along the handle of the hammer when approaching from the side. Unlike the results presented in section 5.3.2, girdle distributions are used to model the orientation of the hand for the side approach. While this is encouraging, the algorithm learned four clusters instead of one. This is most likely due to the spatially distinct hand positions used to grasp the hammer's handle. Also, notice that the algorithm selected a girdle distribution to model grasps near the head of the hammer. This is surprising given that a fixed hand orientation was used to grasp the hammer's head.

The finger configurations used by the teleoperator to grasp the hammer's handle were very different than those used to grasp the hammer's head. When approaching the object from the side, power grasps that maximized the contact surface area between the hand and the handle were more likely to be used. Conversely, precision grasps that mainly used the finger tips were employed when grasping the head of the hammer. These differences in hand shape can be seen in figure 6.3(f). The red ellipsoid represents the finger configurations used to grasp the hammer from above. The large volume of the ellipsoid is due to the exploration strategy employed by the teleoperator: the hand was continually opened and closed on this portion of the object. Hence, there was a large variance in finger configuration. Also, notice that the blue, magenta, green, and orange ellipsoids are spatially distinct from the red eigengrasp cluster.

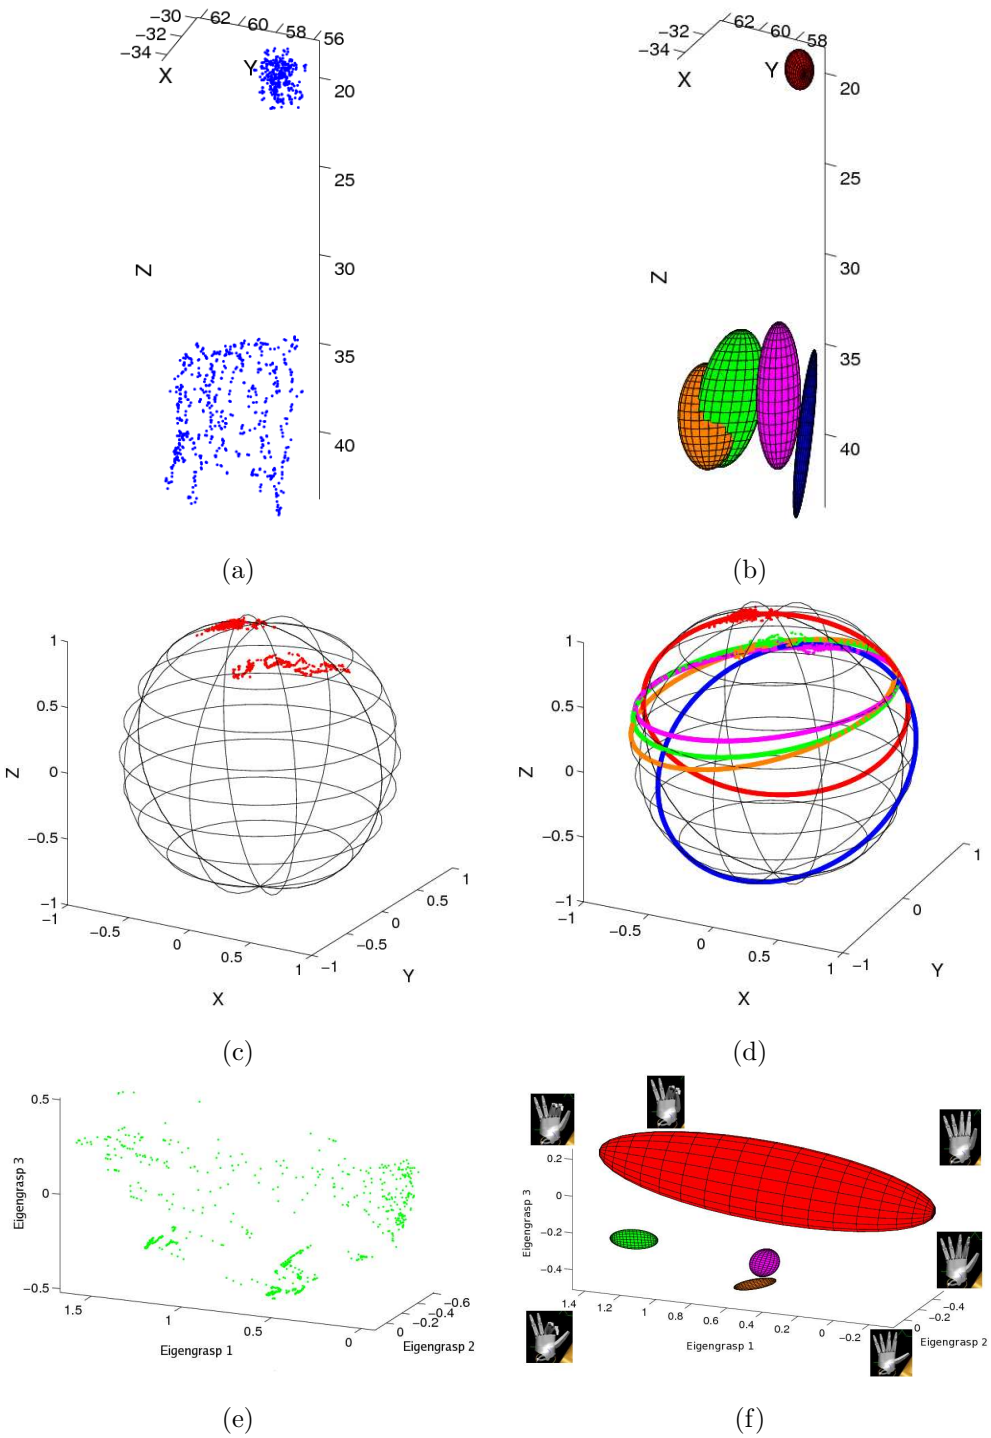

Figure 6.3: The training examples and learned affordance model for the hammer. (a) The position of the hand; (b) The position component of the learned affordance model; (c) The orientation of the hand; (d) The orientation component of the learned affordance model; (e) The finger configuration of the hand; (f) The finger configuration component of the learned affordance model.

6.5 Discussion

In this chapter, our clustering algorithm was extended by incorporating the shape of the hand into the grasp affordance learning process. This was accomplished by constructing PDF's over hand pose and finger configuration. Because a robot's end-effector can have many degrees of freedom, a low dimensional representation of finger configuration is constructed using the hand's eigengrasps. Experimental results demonstrate the ability of our algorithm to learn a small number of canonical pre-shapes that may be used to grasp a given object based on experience provided by a human teleoperator.

One possibility for future work could be the development of an algorithm that identifies grasp events in the experience provided by the teleoperator. This algorithm could be employed as a preprocessing step for the clustering algorithm, and would replace the time consuming method currently being used. This would also be a help in learning more compact finger configuration clusters. For example, the blue and red ellipsoids learned for the bar and the hammer correspond to portions of the object where the teleoperator continuously opened and closed the hand. As a result, some of the finger configurations recorded as grasps actually correspond to times when the fingers were not in contact with the object. This is undesirable in that we wish to model the finger configurations *actually* used to grasp the object.

Chapter 7

Affordance-Based Manipulator Control

7.1 Problem Definition and Algorithm

Once the grasp affordances of an object have been learned, how might an agent use the learned representation as a means of controlling and planning grasping actions? First, the robot determines the identity and pose of the object it intends to grasp. Solving this problem is beyond the scope of this thesis, and it is assumed that the agent has some means of acquiring this information. Next, because the grasp affordances of the object are expressed in an object-centered coordinate frame, the model is aligned with the current pose of the object as expressed in a coordinate frame relative to the robot. Then, the agent selects one of the affordances from the learned model. Because this may be accomplished in a variety of ways (e.g., as a function of task semantics), it is assumed that the affordance of interest to the agent is known *a priori*. Each of the learned clusters in a grasp affordance model provides a set of feasible goals for the robot's hand, but for a given cluster the agent must ultimately limit the possibilities to a single hand pose for the purpose of grasp execution. Our approach to this problem is to choose the pose that is closest to the current configuration of the robot's hand. Because we assume that the position and orientation of the hand are conditionally independent given a cluster, goal selection for each of them may be handled separately. Once a goal pose for the hand has been computed, a trajectory generator is used to execute the desired motion.

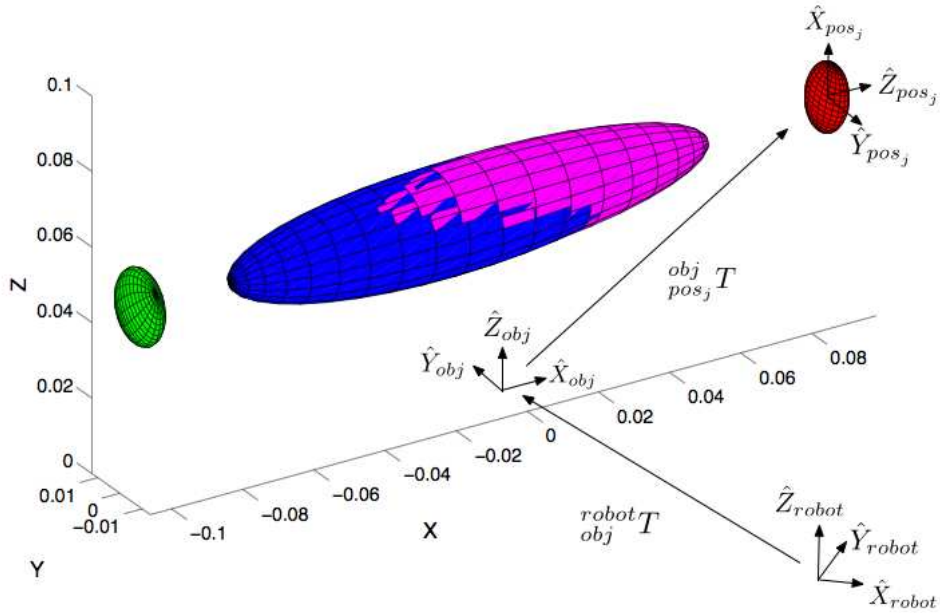

Figure 7.1: A visualization of the ${}^{robot}_{obj}T$ and ${}^{obj}_{pos_j}T$ homogeneous transformations.

7.2 Expressing Grasp Affordances in a Robot-Centered Coordinate Frame

Let ${}^{robot}_{obj}T$ denote the homogeneous transform that specifies the pose of the object relative to the robot. Aligning the corresponding grasp affordance model with the object's current pose is accomplished by applying ${}^{robot}_{obj}T$ to each of the learned clusters in the model. Because each cluster is represented by a parameterized function, the desired transformation is obtained by applying ${}^{robot}_{obj}T$ to the relevant parameters of the position and orientation components of each cluster, respectively.

The position component of cluster j is defined by the parameters $\boldsymbol{\mu}_j$ and $\boldsymbol{\Sigma}_j$, which describe an ellipsoid in $\mathbb{R}^3$. The centroid of this ellipsoid is represented by $\boldsymbol{\mu}_j$, while its orientation is described by the eigenvectors of $\boldsymbol{\Sigma}_j$. Let $\mathbf{V}_j$ denote a matrix that contains the eigenvectors of $\boldsymbol{\Sigma}_j$ as its columns, and $\mathbf{D}_j$ be a diagonal matrix containing the corresponding eigenvalues. Note that since $\mathbf{V}_j$ is an orthonormal matrix, it is also

a proper rotation matrix (Lewis et al. 2006). Hence, the homogeneous transform that expresses the pose of ellipsoid j in terms of the object is defined as follows:

$$ {}^{obj}_{pos_j}T = \begin{bmatrix} \mathbf{V}_j & \boldsymbol{\mu}_j \\ \mathbf{0} & 1 \end{bmatrix}, \tag{7.1} $$

where $\mathbf{0}$ denotes the 1×3 zero vector. Aligning ellipsoid j with the current pose of the object is accomplished by:

$$ {}^{robot}_{pos_j}T = \begin{bmatrix} \hat{\mathbf{V}}_j & \hat{\boldsymbol{\mu}}_j \\ \mathbf{0} & 1 \end{bmatrix} = {}^{robot}_{obj}T \; {}^{obj}_{pos_j}T. \tag{7.2} $$

The resulting ellipsoid's pose relative to the robot is described by $\hat{\boldsymbol{\mu}}_j$ and $\hat{\mathbf{V}}_j$, respectively. Furthermore, the parameters of the corresponding Gaussian distribution are $\hat{\boldsymbol{\mu}}_j$ and $\hat{\mathbf{V}}_j \; \mathbf{D}_j \; \hat{\mathbf{V}}_j^T$. The latter expression is derived using the Singular Value Decomposition (SVD) of symmetric and positive definite matrices (Lewis et al. 2006). Figure 7.1 provides a visualization of the ${}^{robot}_{obj}T$ and ${}^{obj}_{pos_j}T$ transformations.

The orientation component of cluster j is represented by either a Dimroth-Watson or girdle distribution. If a cluster has a Dimroth-Watson distribution as its orientation component, then the parameter of interest is $\mathbf{u}_j$. Let ${}^{obj}_{ort_j}T$ be the homogeneous transform that corresponds to the orientation specified by the unit quaternion $\mathbf{u}_j$. Note that the translation component of ${}^{obj}_{ort_j}T$ is irrelevant because we are dealing with the orientation component of the model, and it is therefore set to the 3×1 zero vector. Aligning ${}^{obj}_{ort_j}T$ with the current pose of the object is accomplished by:

$$ {}^{robot}_{ort_j}T = {}^{robot}_{obj}T \; {}^{obj}_{ort_j}T. \tag{7.3} $$

If the orientation component of cluster j is a girdle distribution the same coordinate frame transformation is applied to both of its parameters $\mathbf{u}_{1j}$ and $\mathbf{u}_{2j}$.

7.3 Extracting Position Goals

Each of the learned position clusters in a grasp affordance model describes a set of goal positions for the robot's hand. The position component of each cluster is represented by a multivariate Gaussian distribution. The set of points in $\mathbb{R}^3$ for which $p(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \geq \beta$ (where β is a threshold parameter that is selected to encompass the first standard deviation) defines the volume of an ellipsoid. Any position within

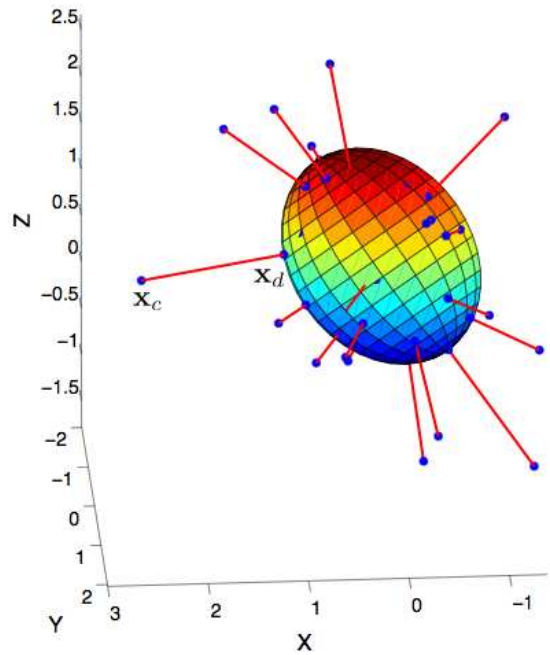

Figure 7.2: A visualization of the position goal selection algorithm.

this ellipsoid is considered to be a feasible grasp point. This set of possibilities is narrowed down to a single goal position for the purposes of control by finding the location on the surface of the ellipsoid that is closest to the current position of the robot's hand (assuming that the hand is not already within the volume). Finding this point is equivalent to solving a nonlinear optimization problem with a constraint. Let $\mathbf{x}_c$ and $\mathbf{x}_d$ denote the current and desired position of the robot's end-effector. The optimization process finds the $\mathbf{x}_d$ that minimizes:

$$E_x = (\mathbf{x}_c - \mathbf{x}_d)^T(\mathbf{x}_c - \mathbf{x}_d), \tag{7.4}$$

subject to the constraint:

$$(\mathbf{x}_d - \boldsymbol{\mu}_j)^T \boldsymbol{\Sigma}_j^{-1}(\mathbf{x}_d - \boldsymbol{\mu}_j) = 1, \tag{7.5}$$

which is the equation of the ellipsoid for cluster j. There are many different ways to solve this type of optimization problem, but the general approach is to reformulate the constrained optimization problem as an unconstrained optimization problem, and to solve a sequence of quadratic programming problems (Gill et al. 1981; Gonzaga et al. 2003).

Figure 7.2 illustrates the result of this process for several different hand positions. Each point that is not on the surface of the ellipsoid corresponds to a particular selection for $\mathbf{x}_c$. Each line that connects two points represents the orthogonal projection of $\mathbf{x}_c$ onto the surface of the ellipsoid, where $\mathbf{x}_d$ is the point on the surface. Note that equation 7.4 corresponds to the squared distance between $\mathbf{x}_c$ and $\mathbf{x}_d$. As a result, the length of each line segment is equal to $\sqrt{E_x}$.

7.4 Extracting Orientation Goals

When selecting goal orientations, different approaches are used depending on whether or not the orientation component of cluster j is a Dimroth-Watson or girdle distribution. Let $\mathbf{q}_c$ and $\mathbf{q}_d$ be the current and desired orientations of the robot's end-effector. In the case of the Dimroth-Watson distribution goal selection is trivial:

$$\mathbf{q}_d = \mathbf{u}_j. \tag{7.6}$$

If the orientation component of cluster j corresponds to a girdle distribution any orientation along the great circle spanned by $\mathbf{u}_{1j}$ and $\mathbf{u}_{2j}$ is viable. The closest orientation is found by projecting the current orientation onto the great circle as follows:

$$\mathbf{q}_d = \frac{proj\left(\mathbf{u}_{1j}, \mathbf{q}_c\right) + proj\left(\mathbf{u}_{2j}, \mathbf{q}_c\right)}{\left|\left|proj\left(\mathbf{u}_{1j}, \mathbf{q}_c\right) + proj\left(\mathbf{u}_{2j}, \mathbf{q}_c\right)\right|\right|}, \tag{7.7}$$

where:

$$proj\left(\mathbf{a}, \mathbf{b}\right) = \frac{\mathbf{a}^T\mathbf{b}}{||\mathbf{a}||}\frac{\mathbf{a}}{||\mathbf{a}||}, \tag{7.8}$$

denotes the projection of the vector $\mathbf{b}$ onto the vector $\mathbf{a}$.

7.5 Executing Grasping Actions

The methods presented in sections 7.3 and 7.4 can be used to provide a reach goal for a controller that drives a robot's hand to a specific pose for grasping. One common method of trajectory generation is called *resolved rate* control (Barker and McKinney 1989; Angeles and Mathur 1990; Craig 2005). . Based on the error vector between the current and desired pose of the robot's end-effector, a resolved rate controller incrementally displaces the manipulator's joints in a way that moves the end-effector

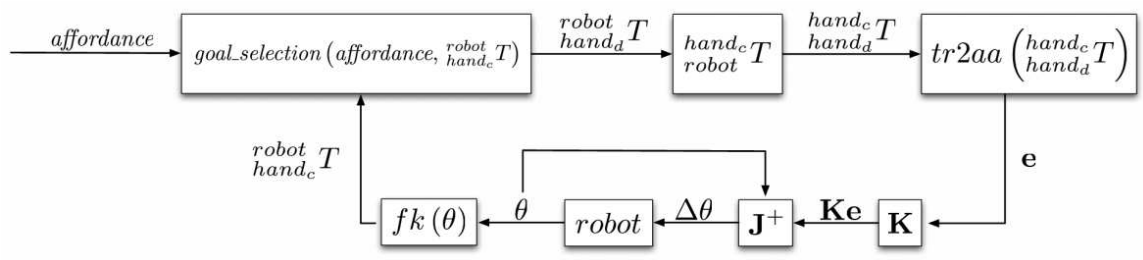

Figure 7.3: A Cartesian controller that makes use of affordances to move a robot's end-effector to a specific pose for grasping.

closer to its goal. During each control cycle, the following control law is used to move the joints of the robot's arm:

$$\Delta\theta = \mathbf{J}^{+}\mathbf{K}\mathbf{e}. \tag{7.9}$$

where:

$$\mathbf{e} = \begin{bmatrix} \mathbf{e}_{pos} \\ \mathbf{e}_{ort} \end{bmatrix}. \tag{7.10}$$

Here, $\Delta\theta$ denotes the relative motions of the arm joints to be executed by the robot, $\mathbf{J}^{+}$ represents the pseudoinverse of the arm Jacobian expressed in a robot centered coordinate frame, $\mathbf{K}$ is a diagonal gain matrix, and $\mathbf{e}$ is the error in end-effector pose. The error vector $\mathbf{e}$ has a position and orientation component, which are denoted by $\mathbf{e}_{pos}$ and $\mathbf{e}_{ort}$, respectively. Note that $\mathbf{e}_{ort}$ is described using an axis angle representation.

To ensure that the hand is always moving towards the closest feasible grasp, the goal is recomputed at the beginning of each control cycle until the controller converges. Figure 7.3 illustrates this process. First, the current pose of the robot's end-effector and the affordance of interest to the agent are supplied to the goal selection algorithm. This generates a homogeneous transform that represents the desired pose for the hand relative to the robot, which is denoted by ${}^{robot}_{hand_d}T$. Then, the error in hand pose is computed $\left({}^{hand_c}_{hand_d}T = {}^{hand_c}_{robot}T\ {}^{robot}_{hand_d}T\right)$. Next, the rotation matrix of ${}^{hand_c}_{hand_d}T$ is converted to an axis angle representation, and appended to the translation component of the transform. At this point, equation 7.9 is applied. The resulting joint displacements are executed by the robot, and a new set of joint angles are reported. The forward kinematic model of the robot's arm is then used to update the pose of the hand, and the entire process repeats until the controller converges.

Chapter 8

Planning and Executing Grasps: An Affordance-Based Approach

8.1 Problem Definition and Algorithm

Chapter 7 presented a controller that utilizes the grasp affordances of an object to execute grasping actions. This chapter provides a preliminary evaluation of this controller using Robonaut. A teleoperator demonstrates some of the feasible grasps for an object, and the learned representation is used to autonomously control the robot.

8.2 Experimental Procedure

After an affordance model has been learned for an object, the robot selects one of the available affordances, and attempts to grasp the object using different starting locations for the hand. Once the controller has converged, the fingers are positioned using the mean of the eigengrasp cluster that is associated with the selected affordance. A grasp is considered to be successful if the object lies within the aperture of the robot's hand. This is determined through visual confirmation by a human.

In addition to the method presented in section 7.3, an alternative approach to selecting goal positions is evaluated. This method projects the current position of the robot's end-effector onto the major axis of the Gaussian ellipsoid. If the resulting position lies outside of the ellipsoid, the goal is selected to be the point at which the ellipsoid's major axis and surface intersect. This approach ensures that two of the three degrees of freedom for the selected goal are equal to those of the ellipsoid's

centroid, which is the maximum likelihood estimate of hand position for a given cluster.

8.3 Experimental Results

Only a single control experiment was performed because time with the robot was very limited. The object used in the experiment is the bar presented in chapters 5 and 6. In this case, the teleoperator only demonstrates the overhand and underhand grasp configurations. Figure 8.1 shows the collected samples, as well as the learned affordance model for the bar. Notice that both sets of grasping options were discovered by the clustering algorithm. Furthermore, the low dimensional eigengrasp clusters overlap considerably because the finger configurations used for both grasp approaches are very similar. The finger configurations shown in figure 8.1(f) correspond to the centroids of the red and blue ellipsoids. These are the hand shapes that were used for each grasp attempted by the robot.

When using the learned affordance model for control, the robot successfully executed an underhanded grasp on four out of the five attempts. These successes were achieved by projecting the current position of the hand onto the major axis of the red ellipsoid shown in figure 8.1(b). However, when the current position of the hand was projected onto the ellipsoid's surface, the hand ended up in a position below the bar that prevented the robot from performing a successful grasp.

The robot was unable to grasp the bar using an overhanded grasp. The controller consistently reached past the bar along the positive x axis (which points away from the robot). One possible explanation for this is that the object moved during the data collection process. Evidence to support this hypothesis is shown in figure 8.2. Notice that there is a significant difference in x between the two ellipsoids (approximately 8 centimeters from centroid to centroid). Previous experiments with the bar show a smaller distance between these two ellipsoids in the x direction (approximately 4 centimeters). Furthermore, in the previous experiments the bar was rotated about its major axis toward the robot. Hence, some separation in x was to be expected. However, the bar was not rotated in this experiment, and one would expect the difference to be smaller.

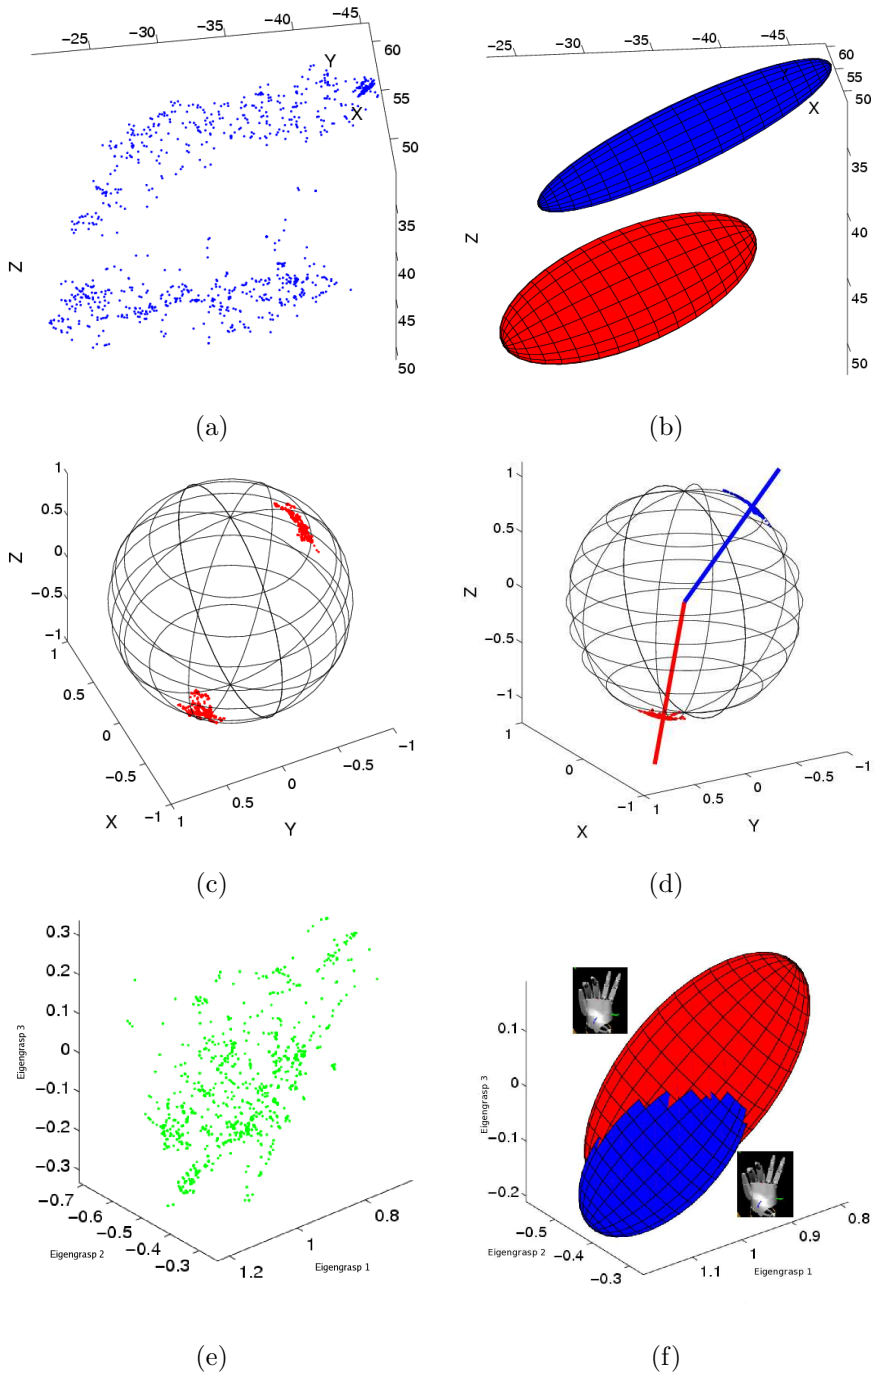

Figure 8.1: The training examples and learned affordance model for the bar. (a) The position of the hand; (b) The position component of the learned affordance model; (c) The orientation of the hand; (d) The orientation component of the learned affordance model; (e) The finger configuration of the hand; (f) The finger configuration component of the learned affordance model, where the hand shapes shown correspond to the centroids of each ellipsoid.

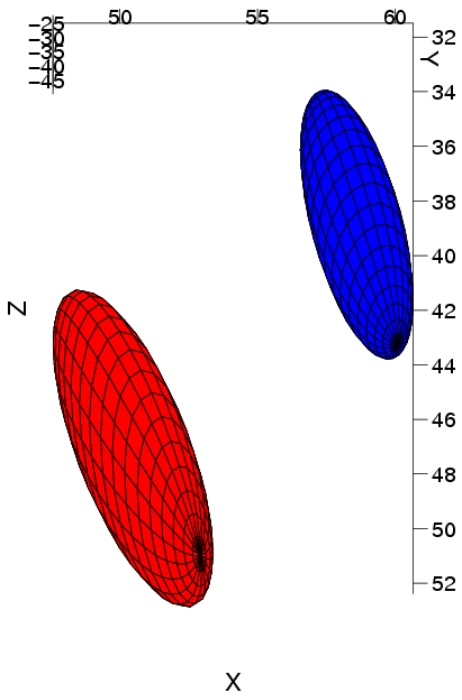

Figure 8.2: A side view of the position clusters learned for the bar.

Regardless of whether or not the robot was able to achieve a hand configuration that placed the object within the aperture of the hand, there were many encouraging aspects of the experiment. For example, the controller consistently drove the hand to an orientation that rendered the surface of the palm approximately parallel to the top and bottom faces of the bar (depending on which grasp was attempted). Also, the hand shape selected for each grasp was appropriate given the bar's geometry.

8.4 Discussion

In this chapter, an affordance-based grasp controller was evaluated with respect to each of the learned affordances for a given object. Preliminary results suggest that our approach is capable of successfully grasping a metallic bar. This work represents a key step toward successfully using the grasp affordance representation presented in this thesis to plan and execute grasping actions. However, additional experiments are needed to more formally assess the quality of the controller. Furthermore, by addressing some of the perception problems associated with grasping (e.g., identifying

the pose of the object relative to the robot), one would not have to worry about the object being moved during training or execution.

Another possibility for future work includes developing a more robust approach to generating goal positions for the hand when grasping. One way of accomplishing this could be to generate goals based on the eccentricity of a given ellipsoid. For example, if the ellipsoid is nearly spherical, then the centroid should be used as the grasping position. However, if the ellipsoid has one principal axis, the method presented in this chapter might be more appropriate. Another alternative is to investigate other distributions for modeling the position of the hand when grasping.

Chapter 9

Conclusions and Future Work

9.1 Conclusions

In this thesis, I present an approach that learns the grasp affordances of objects through human demonstration and uses the learned representation to plan and execute grasping actions. The major contributions of this work are as follows:

1) A mathematical formulation for representing grasp affordances, and a computational framework for learning them based on human demonstration and robot teleoperation. Each canonical grasp is represented by a joint probability distribution over the pose of the hand and the configuration of the fingers. A mixture model-based approach is employed to describe objects for which multiple classes of grasps exist. One useful property of the proposed representation is that it provides a small number of discrete grasping options. By limiting the branching factor, this could enable an agent to efficiently learn which available grasps are most relevant to performing a particular task.

2) An algorithm for extracting reach goals from the proposed grasp affordance representation. Each affordance represents a set of feasible grasps that one expects to be successful if executed. Given an affordance, goal extraction is accomplished by narrowing the set of feasible grasps down to the one that is nearest to the current pose of the robot's hand.

3) A grasp controller that drives the hand of a robot to a specific position and orientation relative to an object, and positions the fingers for grasping.

This thesis has shown the feasibility of extracting a small number of canonical grasps for an object from a human demonstration. Furthermore, preliminary experiments with a humanoid robot suggest the viability of using the proposed representation to autonomously execute grasps.

9.2 Future Work

In this thesis, I have demonstrated that grasping experience can be derived from observation of a human teacher, or from robot teleoperation. However, it may also come from a robot that is performing a grasping task automatically. For example, Coelho and Grupen (1997) use a haptically-guided approach to find a set of finger contact locations that minimize the net force and torque exerted upon an object. The resulting set of finger configurations, and the corresponding poses for the hand, could then be used to learn the grasp affordances of the object.

Currently, the grasp affordance learning algorithm is limited to the class of objects that can be grasped with a single hand, but many manipulation tasks may require the use of two-handed grasps, the execution of sequences of re-grasps, or the cooperation of multiple agents. For example, suppose an agent is handed a cup. If the top grasp is already being used, the agent must explicitly acknowledge this and use a different grasp. One possible solution to this problem is to define joint probability distributions over the position, orientation, and finger configuration of both hands.

One could also explore the use of affordances for higher level planning and learning. The affordance representation presented in this thesis captures the syntax of grasping (i.e., what grasps are possible for a given object), but it does not take into account the semantics of grasping (how an object is to be used in the larger context of a task). This distinction, which is drawn by Gibson, is a critical one for a learning agent. When a task is presented, the syntax of interacting with a specific object can be readily accessed and used, but the learning agent is then left with the problem of determining a sequence of actions that will achieve its goal.

Finally, one must address the perception problems that arise when taking an indirect approach to learning grasp affordances. Ultimately, an agent must establish a mapping between perceivable features to parameterized grasping actions through some sort of object-centered representation. Wang (2007) has recently presented an

approach that recognizes the identity and pose of objects based on visual features. By using this intermediate object representation one can establish an indirect connection between visual features and the hand preshape clusters learned by my algorithm. This approach could also allow one to learn representations that are not specific to any particular object, but to components of objects. Thus, if a novel object is composed of parts similar to those in previous experience, the robot should still be able to grasp the object.

Bibliography

Ambrose, R. O., Aldridge, H., Askew, R. S., Burridge, R. R., Bluethmann, W., Diftler, M., Lovchik, C., Magruder, D. and Rehnmark, F. (2000), 'Robonaut: Nasa's space humanoid', *IEEE Intelligent Systems* **15**(4), 57–63.

Angeles, J. and Mathur, S. (1990), Resolved-rate control of redundant manipulators with elimination of non-conservation effects, *in* 'The Fifth International Symposium on Robotics research', MIT Press, Cambridge, MA, USA, pp. 209–216.

Barker, L. K. and McKinney, Jr., W. S. (1989), Optimized resolved rate control of seven-degree-of-freedom laboratory telerobotic manipulator (ltm) with application to three-dimensional graphics simulation, Technical Report 2938, NASA/Langley Research Center.

Bekey, G. A., Liu, H., Tomovic, R. and Karplus, W. J. (1993), 'Knowledge-based control of grasping in robot hands using heuristics from human motor skills', *IEEE Transactions on Robotics and Automation* **9**(6), 709–722.

Beyer, K., Goldstein, J., Ramakrishnan, R. and Shaft, U. (1999), 'When is "nearest neighbor" meaningful?', *Lecture Notes in Computer Science* **1540**, 217–235.

Biernacki, C., Celeux, G. and Govaert, G. (2000), 'Assessing a mixture model for clustering with the integrated completed likelihood', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(7), 719–725.

Brock, O., Fagg, A. H., Grupen, R. A., Karuppiah, D., Platt, R. and Rosenstein, M. (2005), 'A framework for humanoid control and intelligence', *International Journal of Humanoid Robotics* **2**(3), 301–336.

Ciocarlie, M., Goldfeder, C. and Allen, P. (2007), Dexterous grasping via eigengrasps: A low-dimensional approach to a high-complexity problem, *in* 'Proceedings of the Robotics: Science & Systems 2007 Workshop - Sensing and Adapting to the Real World'. Electronically published.

Coelho, Jr., J. A. and Grupen, R. A. (1997), 'A control basis for learning multifingered grasps', *Journal of Robotic Systems* **14**(7), 545–557.

Coelho, Jr., J. A., Piater, J. and Grupen, R. A. (2000), 'Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot', *Robotics*

and Autonomous Systems Journal, special issue on Humanoid Robots **37**(2–3), 195–219.

Craig, J. J. (2005), *Introduction to Robotics: Mechanics and Control*, 3rd edn, Pearson Prentice Hall, Upper Saddle River, NJ.

de Granville, C., Southerland, J. and Fagg, A. H. (2006), Learning grasp affordances through human demonstration, *in* 'Proceedings of the International Conference on Development and Learning'. electronically published.

Dempster, A., Laird, N. and Rubin, D. (1977), 'Maximum likelihood estimation from incomplete data via the EM algorithm', *Journal of the Royal Statistical Society, Series B* **39**(1), 1–38.

Fagg, A. H., Rosenstein, M. T., Platt, Jr., R. and Grupen, R. A. (2004), Extracting user intent in mixed initiative teleoperator control, *in* 'Proceedings of the American Institute of Aeronautics and Astronautics Intelligent Systems Technical Conference'. Electronically published.

Gibson, J. J. (1966), *The Senses Considered as Perceptual Systems*, Allen and Unwin.

Gibson, J. J. (1977), The theory of affordances, *in* R. E. Shaw and J. Bransford, eds, 'Perceiving, Acting, and Knowing', Lawrence Erlbaum, Hillsdale.

Gill, P. E., Murray, W. and Wright, M. H. (1981), *Practical optimization*, London: Academic Press, 1981.

Gonzaga, C. C., Karas, E. and Vanti, M. (2003), 'A globally convergent filter method for nonlinear programming', *SIAM Journal on Optimization* **14**(3), 646–669.

Hand, D. J., Smyth, P. and Mannila, H. (2001), *Principles of Data Mining*, MIT Press, Cambridge, MA, USA.

Lewis, J. M., Lakshmivarahan, S. and Dhall, S. (2006), *Dynamic Data Assimilation: A Least Squares Approach*, Cambridge University Press, NY, USA.

Mardia, K. V. and Jupp, P. E. (1999), *Directional Statistics*, Wiley Series in Probability and Statistics, Wiley, Chichester, West Sussex, England.

Miller, A., Knoop, S., Christensen, H. and Allen, P. (2003), Automatic grasp planning using shape primitives, *in* 'Proceedings of the IEEE International Conference on Robotics and Automation', pp. 1824–2829.

Montesano, L., Lopes, M., Bernardino, A. and Santos-Victor, J. (2007), Modeling affordances using Bayesian networks, *in* 'IEEE/RSJ International Conference on Intelligent Robots and Systems', pp. 4102–4107.

Piater, J. and Grupen, R. (2002), Learning appearance features to support robotic manipulation, *in* 'Proceedings of the Cognitive Vision Workshop'. Electronically published.

Platt, Jr., R., Fagg, A. H. and Grupen, R. A. (2002), Nullspace composition of control laws for grasping, *in* 'Proceedings of the International Conference on Intelligent Robots and Systems', pp. 1717–1723.

Rancourt, D., Rivest, L.-P. and Asselin, J. (2000), 'Using orientation statistics to investigate variations in human kinematics', *Applied Statistics* **49**(1), 81–94.

Rivest, L.-P. (2001), 'A directional model fo the statistical analysis of movement in three dimensions', *Biometrika* **88**(3), 779–791.

Santello, M., Flanders, M. and Soechting, J. F. (1998), 'Postural hand synergies for tool use', *Journal of Neuroscience* **18**(23), 10105–10115.

Stoytchev, A. (2005), Toward learning the binding affordances of objects: A behavior-grounded approach, *in* 'Proceedings of the AAAI Spring Symposium on Developmental Robotics'. Electronically published.

Sweeney, J. D. and Grupen, R. (2007), A model of shared grasp affordances from demonstration, *in* 'IEEE-RAS International Conference on Humanoid Robots'. Session FRa2, electronically published.

Wang, D. (2007), A 3D feature-based object recognition system for grasping, Master's thesis, School of Computer Science, University of Oklahoma, Norman, OK.

Appendix A

A.1 Normalization Terms

Let $\mathbf{q} \in \mathbb{R}^4$ represent a quaternion. When deriving the normalization terms of the Dimroth-Watson and girdle distributions, it is convenient to describe the surface of a hypersphere using a spherical coordinate system that is defined by three angles (θ, ξ_1, and ξ_2), and a radius (r). One possible transformation between spherical coordinates and quaternions is:

$$\mathbf{q}\left(\theta, \xi_1, \xi_2, r\right) = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} = \begin{bmatrix} r\cos\xi_1 \sin\frac{\theta}{2} \\ r\sin\xi_1 \sin\frac{\theta}{2} \\ r\cos\xi_2 \cos\frac{\theta}{2} \\ r\sin\xi_2 \cos\frac{\theta}{2} \end{bmatrix}, \tag{A.1}$$

where $0 \leq \theta \leq \pi$, $0 \leq \xi_1 < 2\pi$, and $0 \leq \xi_2 < 2\pi$. The derivative of $\mathbf{q}\left(\right)$ with respect to θ, ξ_1, ξ_2, and r results in a matrix of partial derivatives called the Jacobian:

$$\begin{aligned} \mathbf{J}\left(\theta, \xi_1, \xi_2, r\right) &= \begin{bmatrix} \frac{\partial q_1}{\partial \theta} & \frac{\partial q_1}{\partial \xi_1} & \frac{\partial q_1}{\partial \xi_2} & \frac{\partial q_1}{\partial r} \\ \frac{\partial q_2}{\partial \theta} & \frac{\partial q_2}{\partial \xi_1} & \frac{\partial q_2}{\partial \xi_2} & \frac{\partial q_2}{\partial r} \\ \frac{\partial q_3}{\partial \theta} & \frac{\partial q_3}{\partial \xi_1} & \frac{\partial q_3}{\partial \xi_2} & \frac{\partial q_3}{\partial r} \\ \frac{\partial q_4}{\partial \theta} & \frac{\partial q_4}{\partial \xi_1} & \frac{\partial q_4}{\partial \xi_2} & \frac{\partial q_4}{\partial r} \end{bmatrix} \\ &= \begin{bmatrix} \frac{1}{2} r\cos\xi_1 \cos\frac{\theta}{2} & -r\sin\xi_1 \sin\frac{\theta}{2} & 0 & \cos\xi_1 \sin\frac{\theta}{2} \\ \frac{1}{2} r\sin\xi_1 \cos\frac{\theta}{2} & r\cos\xi_1 \sin\frac{\theta}{2} & 0 & \sin\xi_1 \sin\frac{\theta}{2} \\ -\frac{1}{2} r\cos\xi_2 \sin\frac{\theta}{2} & 0 & -r\sin\xi_2 \cos\frac{\theta}{2} & \cos\xi_2 \cos\frac{\theta}{2} \\ -\frac{1}{2} r\sin\xi_2 \sin\frac{\theta}{2} & 0 & r\cos\xi_2 \cos\frac{\theta}{2} & \sin\xi_2 \cos\frac{\theta}{2} \end{bmatrix} \end{aligned}$$

To integrate functions involving $\mathbf{q}\,()$ over S3, the determinant of the Jacobian is needed:

$$
\begin{aligned}
\det \mathbf{J}\left(\theta, \xi_1, \xi_2, r\right) &= \frac{1}{2} r^2 \cos\xi_1 \cos\frac{\theta}{2} \cos\xi_1 \sin\frac{\theta}{2} \left[-r \sin\xi_2 \cos\frac{\theta}{2} \sin\xi_2 \cos\frac{\theta}{2} \right. \\
&\quad \left. -r \cos\xi_2 \cos\frac{\theta}{2} \cos\xi_2 \cos\frac{\theta}{2} \right] \\
&\quad + r \sin\xi_1 \sin\frac{\theta}{2} \left\{ \frac{1}{2} r \sin\xi_1 \cos\frac{\theta}{2} \left[-r \sin\xi_2 \cos\frac{\theta}{2} \sin\xi_2 \cos\frac{\theta}{2} \right. \right. \\
&\quad \left. -r \cos\xi_2 \cos\frac{\theta}{2} \cos\xi_2 \cos\frac{\theta}{2} \right] \\
&\quad + \frac{1}{2} \sin\xi_1 \sin\frac{\theta}{2} \left[-r^2 \cos\xi_2 \sin\frac{\theta}{2} \cos\xi_2 \cos\frac{\theta}{2} \right. \\
&\quad \left. \left. -r^2 \sin\xi_2 \sin\frac{\theta}{2} \sin\xi_2 \cos\frac{\theta}{2} \right] \right\} \\
&\quad + r \cos\xi_1 \sin\frac{\theta}{2} \cos\xi_1 \sin\frac{\theta}{2} \left[-r^2 \frac{1}{2} \cos\xi_2 \sin\frac{\theta}{2} \cos\xi_2 \cos\frac{\theta}{2} \right. \\
&\quad \left. -r^2 \frac{1}{2} \sin\xi_2 \sin\frac{\theta}{2} \sin\xi_2 \cos\frac{\theta}{2} \right] \\
&= \frac{1}{2} r^3 \cos^2\xi_1 \cos^3\frac{\theta}{2} \sin\frac{\theta}{2} \left[-\sin^2\xi_2 - \cos^2\xi_2 \right] \\
&\quad + r^3 \sin\xi_1 \sin\frac{\theta}{2} \left\{ \frac{1}{2} \sin\xi_1 \cos^3\frac{\theta}{2} \left[-\sin^2\xi_2 - \cos^2\xi_2 \right] \right. \\
&\quad \left. + \frac{1}{2} \sin\xi_1 \sin^2\frac{\theta}{2} \cos\frac{\theta}{2} \left[-\cos^2\xi_2 - \sin^2\xi_2 \right] \right\} \\
&\quad + \frac{1}{2} r^3 \cos^2\xi_1 \sin^3\frac{\theta}{2} \cos\frac{\theta}{2} \left[-\cos^2\xi_2 - \sin^2\xi_2 \right] \\
&= -\frac{1}{2} r^3 \left[\cos^2\xi_1 \cos^3\frac{\theta}{2} \sin\frac{\theta}{2} + \sin^2\xi_1 \sin\frac{\theta}{2} \cos^3\frac{\theta}{2} \right. \\
&\quad \left. + \sin^2\xi_1 \sin^3\frac{\theta}{2} \cos\frac{\theta}{2} + \cos^2\xi_1 \sin^3\frac{\theta}{2} \cos\frac{\theta}{2} \right].
\end{aligned}
$$

Note that since $2\cos\frac{\theta}{2}\sin\frac{\theta}{2} = \sin\theta$, the above equation becomes:

$$\begin{aligned}
\det \mathbf{J}\left(\theta, \xi_1, \xi_2, r\right) &= -\frac{1}{2}r^3\left[\frac{1}{2}\cos^2\xi_1\cos^2\frac{\theta}{2}\sin\theta + \frac{1}{2}\sin^2\xi_1\cos^2\frac{\theta}{2}\sin\theta\right. \\
&\quad \left. + \frac{1}{2}\sin^2\xi_1\sin^2\frac{\theta}{2}\sin\theta + \frac{1}{2}\cos^2\xi_1\sin^2\frac{\theta}{2}\sin\theta\right] \\
&= -\frac{1}{4}r^3\sin\theta\left[\cos^2\xi_1\cos^2\frac{\theta}{2} + \sin^2\xi_1\cos^2\frac{\theta}{2}\right. \\
&\quad \left. + \sin^2\xi_1\sin^2\frac{\theta}{2} + \cos^2\xi_1\sin^2\frac{\theta}{2}\right] \\
&= -\frac{1}{4}r^3\sin\theta\left[\cos^2\frac{\theta}{2}\left(\cos^2\xi_1 + \sin^2\xi_1\right)\right. \\
&\quad \left. + \sin^2\frac{\theta}{2}\left(\sin^2\xi_1 + \cos^2\xi_1\right)\right] \\
&= -\frac{1}{4}r^3\sin\theta.
\end{aligned}$$

Because we are concerned with the unit hypersphere $r = 1$, and $\det \mathbf{J}\left(\theta, \xi_1, \xi_2, r\right) = -\frac{1}{4}\sin\theta$.

In order for the Dimroth-Watson distribution to be considered a proper probability density function, the following must be true:

$$\int_S f\left(\mathbf{q}|\mathbf{u}, k\right) d^3S = 1, \tag{A.2}$$

Substituting equation 3.1 into equation A.2 we have:

$$\begin{aligned}
1 &= \int_S F\left(k\right) e^{k\left(\mathbf{q}^T\mathbf{u}\right)^2} d^3S \\
\frac{1}{F\left(k\right)} &= \int_S e^{k\left(\mathbf{q}^T\mathbf{u}\right)^2} d^3S \\
&= \int_{-1}^{1}\int_{-\sqrt{1-q_1^2}}^{\sqrt{1-q_1^2}}\int_{-\sqrt{1-q_1^2-q_2^2}}^{\sqrt{1-q_1^2-q_2^2}}\left(e^{k\left(\mathbf{q}_+^T\mathbf{u}\right)^2} + e^{k\left(\mathbf{q}_-^T\mathbf{u}\right)^2}\right) dq_3 dq_2 dq_1, \qquad \text{(A.3)}
\end{aligned}$$

where:

$$\mathbf{q}_+ = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ +\sqrt{1 - q_1^2 - q_2^2 - q_3^2} \end{bmatrix},$$

and

$$\mathbf{q}_- = \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ -\sqrt{1 - q_1^2 - q_2^2 - q_3^2} \end{bmatrix}.$$

Here, $\mathbf{q}_+$ and $\mathbf{q}_-$ define points on opposite hemispheres. Transforming equation A.3 to spherical coordinates results in:

$$\frac{1}{F(k)} = \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} e^{k\left(\mathbf{q}(\theta,\xi_1,\xi_2,1)^T\mathbf{u}\right)^2} \left|\det \mathbf{J}(\theta,\xi_1,\xi_2,1)\right| d\xi_2 d\xi_1 d\theta, \tag{A.4}$$

where the bounds on the integrals ensure a covering of the entire hypersphere.

Since $\mathbf{u}$ does not influence $F(k)$, we can arbitrarily select it to be $\begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}^T$. Hence, equation A.4 becomes:

$$\begin{aligned}
\frac{1}{F(k)} &= \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} \frac{1}{4}\sin\theta e^{k\cos^2\xi_1\sin^2\frac{\theta}{2}} d\xi_2 d\xi_1 d\theta \\
&= \int_0^{\pi}\int_0^{2\pi} \frac{2\pi}{4}\sin\theta e^{k\cos^2\xi_1\sin^2\frac{\theta}{2}} d\xi_1 d\theta \\
&= \int_0^{\pi}\int_0^{2\pi} \frac{\pi}{2}\sin\theta e^{k\cos^2\xi_1\sin^2\frac{\theta}{2}} d\xi_1 d\theta.
\end{aligned}$$

Solving for $F(k)$ we have:

$$F(k) = \frac{2}{\pi\int_0^{\pi}\int_0^{2\pi}\sin\theta e^{k\cos^2\xi_1\sin^2\frac{\theta}{2}} d\xi_1\, d\theta}. \tag{A.5}$$

The process for deriving the normalization term of the girdle distribution is very similar to the one for the Dimroth-Watson distribution. First, we rewrite A.4 as:

$$\frac{1}{\bar{F}(k)} = \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} e^{k\left[\left(\mathbf{q}(\theta,\xi_1,\xi_2,1)^T\mathbf{u}_1\right)^2+\left(\mathbf{q}(\theta,\xi_1,\xi_2,1)^T\mathbf{u}_2\right)^2\right]} \left|\det \mathbf{J}(\theta,\xi_1,\xi_2,1)\right| d\xi_2 d\xi_1 d\theta \tag{A.6}$$

The choices for $\mathbf{u}_1$ and $\mathbf{u}_2$ are arbitrary as long as they are orthogonal. We select them to be $\begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}^T$ and $\begin{bmatrix} 0 & 1 & 0 & 0 \end{bmatrix}^T$, respectively. As a result, equation A.6 becomes:

$$\begin{aligned}
\frac{1}{\bar{F}(k)} &= \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} \frac{1}{4}\sin\theta e^{k\left[\cos^2\xi_1\sin^2\frac{\theta}{2}+\sin^2\xi_1\sin^2\frac{\theta}{2}\right]} d\xi_2 d\xi_1 d\theta \\
&= \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} \frac{1}{4}\sin\theta e^{k\sin^2\frac{\theta}{2}\left[\cos^2\xi_1+\sin^2\xi_1\right]} d\xi_2 d\xi_1 d\theta \\
&= \int_0^{\pi}\int_0^{2\pi}\int_0^{2\pi} \frac{1}{4}\sin\theta e^{k\sin^2\frac{\theta}{2}} d\xi_2 d\xi_1 d\theta \\
&= \int_0^{\pi}\int_0^{2\pi} \frac{2\pi}{4}\sin\theta e^{k\sin^2\frac{\theta}{2}} d\xi_1 d\theta \\
&= \int_0^{\pi} \frac{4\pi^2}{4}\sin\theta e^{k\sin^2\frac{\theta}{2}} d\theta \\
&= \int_0^{\pi} \pi^2\sin\theta e^{k\sin^2\frac{\theta}{2}} d\theta \\
&= \int_0^{\pi} \pi^2\sin\theta e^{\frac{k}{2}(1-\cos\theta)} d\theta.
\end{aligned} \tag{A.7}$$

This last step is achieved by noticing that $\sin^2\frac{\theta}{2} = \frac{1}{2}(1-\cos\theta)$. Now, if we let $v = 1-\cos\theta$, and $dv = \sin\theta\, d\theta$, we can substitute these into equation A.7, and solve for $\bar{F}(k)$ as follows:

$$\begin{aligned}\bar{F}(k) &= \frac{1}{\int_0^2 \pi^2 e^{\frac{k}{2}v} dv} \\ &= \frac{k}{2\pi^2\left(e^k-1\right)}.\end{aligned} \tag{A.8}$$

A.2 Maximum Likelihood Estimates of the Concentration Parameters

The derivation of equation 3.12 begins with the weighted log-likelihood function for the Dimroth-Watson distribution. Here, j is the index of a particular cluster, and α_{ji} denotes the probability that sample i belongs to cluster j (with the constraint that $\forall i$: $\sum_{j=1}^{M}\alpha_{ji} = 1$). The weighted log-likelihood function may now be defined as:

$$\begin{aligned}WLL &= \log\left(\prod_{i=1}^{N} f\left(\mathbf{q}_i|\mathbf{u}_j, k_j\right)^{\alpha_{ji}}\right) \\ &= \sum_{i=1}^{N}\log\left(f\left(\mathbf{q}_i|\mathbf{u}_j, k_j\right)^{\alpha_{ji}}\right) \\ &= \sum_{i=1}^{N}\alpha_{ji}\log f\left(\mathbf{q}_i|\mathbf{u}_j, k_j\right) \\ &= \sum_{i=1}^{N}\alpha_{ji}\left(\log F\left(k_j\right) + \log e^{k_j\left(\mathbf{q}_i^T\mathbf{u}_j\right)^2}\right) \\ &= \sum_{i=1}^{N}\alpha_{ji}\log F\left(k_j\right) + \sum_{i=1}^{N}\alpha_{ji}k_j\left(\mathbf{q}_i^T\mathbf{u}_j\right)^2 \\ &= \log F\left(k_j\right)\sum_{i=1}^{N}\alpha_{ji} + k_j\sum_{i=1}^{N}\alpha_{ji}\left(\mathbf{q}_i^T\mathbf{u}_j\right)^2.\end{aligned}$$

Taking the derivative of WLL with respect to k_j, and setting the result to zero results in:

$$\begin{aligned}
\frac{\partial WLL}{\partial k_j} &= \frac{F'(k_j)}{F(k_j)} \sum_{i=1}^{N} \alpha_{ji} + \sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}_j\right)^2 \\
0 &= \frac{F'(k_j)}{F(k_j)} \sum_{i=1}^{N} \alpha_{ji} + \sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}_j\right)^2 \\
-\frac{F'(k_j)}{F(k_j)} \sum_{i=1}^{N} \alpha_{ji} &= \sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}_j\right)^2 \\
G(k_j) \equiv \frac{F'(k_j)}{F(k_j)} &= -\frac{\sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}_j\right)^2}{\sum_{i=1}^{N} \alpha_{ji}}.
\end{aligned}$$

Note that if there is only a single cluster, then $\forall i \ \alpha_{1i} = 1$, and the above equation reduces to:

$$G(k) \equiv \frac{F'(k)}{F(k)} = -\frac{\sum_{i=1}^{N} \left(\mathbf{q}_i^T \mathbf{u}\right)^2}{N},$$

which is equation 3.4.

In the case of a girdle distribution, equations 3.5 and 3.13 may be derived through a similar process.

For computational efficiency, the maximum likelihood estimates of k_j for the Dimroth-Watson and girdle distributions are approximated by:

$$k_j \approx G^{-1}(z_j) = 1.9090 + 3.4599 \log(z_j) \log(1.6376\, z_j), \text{ and}$$

$$k_j \approx \bar{G}^{-1}(\bar{z}_j) = 4.2648 + 4.0254 \log(\bar{z}_j) \log(5.2532\, \bar{z}_j),$$

where

$$z_j = -\frac{\sum_{i=1}^{N} \alpha_{ji} \left(\mathbf{q}_i^T \mathbf{u}_j\right)^2}{\sum_{i=1}^{N} \alpha_{ji}}, \text{ and}$$

$$\bar{z}_j = -\frac{\sum_{i=1}^{N} \alpha_{ji} \left[\left(\mathbf{q}_i^T \mathbf{u}_{1j}\right)^2 + \left(\mathbf{q}_i^T \mathbf{u}_{2j}\right)^2\right]}{\sum_{i=1}^{N} \alpha_{ji}}.$$

Again, if there is only a single cluster, then z_j and $\bar{z}_j$ reduce to:

$$z = -\frac{\sum_{i=1}^{N} \left(\mathbf{q}_i^T \mathbf{u}\right)^2}{N}, \text{ and}$$

$$\bar{z} = -\frac{\sum_{i=1}^{N} \left[\left(\mathbf{q}_i^T \mathbf{u}_1\right)^2 + \left(\mathbf{q}_i^T \mathbf{u}_2\right)^2\right]}{N}.$$