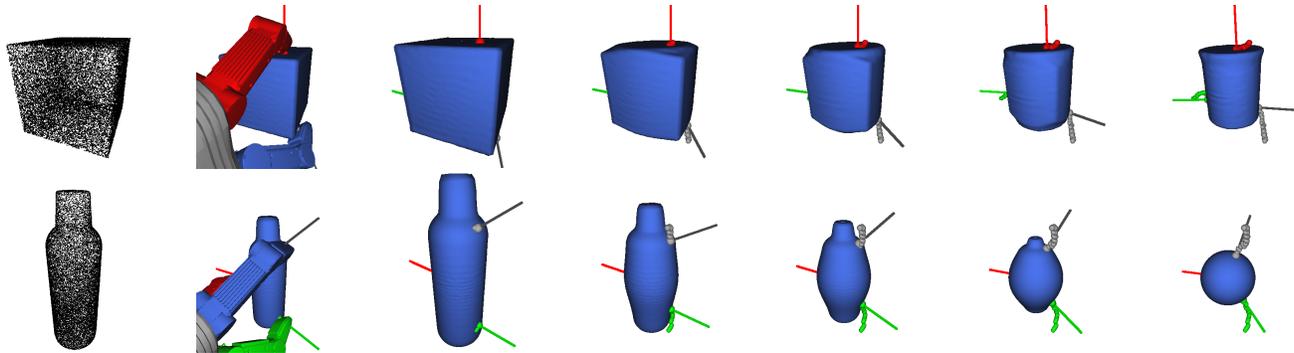


Grasp Moduli Spaces and Spherical Harmonics

Florian T. Pokorny

Yasemin Bekiroglu

Danica Kragic



Abstract—In this work, we present a novel representation which enables a robot to reason about, transfer and optimize grasps on various objects by representing objects and grasps on them jointly in a common space. In our approach, objects are parametrized using smooth differentiable functions which are obtained from point cloud data via a spectral analysis. We show how, starting with point cloud data of various objects, one can utilize this space consisting of grasps and smooth surfaces in order to continuously deform various surface/grasp configurations with the goal of synthesizing force closed grasps on novel objects. We illustrate the resulting shape space for a collection of real world objects using multidimensional scaling and show that our formulation naturally enables us to use gradient ascent approaches to optimize and simultaneously deform a grasp from a known object towards a novel object.

I. INTRODUCTION

The problem of determining a *suitable representation* which enables a robot to reason about the manipulation of objects in its environment is fundamental to robotics. While most current grasp synthesis approaches are based on non-smooth object representations such as triangle meshes, this work considers *smooth representations of objects that can be continuously varied*. Furthermore, while previously objects such as boxes, cylinders, spheres and quadric surfaces have been considered as primitives for the purpose of grasping and manipulation (see [1] and references therein), we shall argue that shape primitives defined in terms of a spectral analysis of the underlying objects provides a useful alternative. Our work extends the approach of [2], where an infinite dimensional space – the Grasp Moduli Space – for contact configurations on surfaces parametrized by a cylindrical coordinate chart was introduced. It was shown that shape deformation based approaches in this space can

The authors are with the Computer Vision and Active Perception Lab, Centre for Autonomous Systems, School of Computer Science and Communication, KTH Royal Institute of Technology, Stockholm, Sweden, {fpokorny, yaseminb, dani}@kth.se. This work was supported by the EU grants FLEXBOT (FP7-ERC-279933), eSMCs (FP7-ICT-270212) and RoboHow.Cog (FP7-ICT-288533).

be used to synthesize and transfer stable grasps if a smooth surface representation of the object is known. In this work, we utilize a spectral object representation, which has also recently been used in neuroanatomy [3], and show how such a representation can form a basis for a deformation based grasp synthesis framework. Our present work extends the approach of [2] since it allows us to use point cloud data directly. Furthermore, grasps on any genus zero (*i.e.* without holes) surface can in principle be modeled with the techniques we consider here, vastly extending the set of objects usable within this approach. The main contributions of this paper can be summarized as follows:

- a) We discuss the infinite-dimensional shape spaces \mathcal{M}^{sph} and \mathcal{M}^{rad} for representing shapes for the purpose of robotic grasping, and we investigate convexity properties of \mathcal{M}^{rad} in particular.
- b) For a set of example surfaces, we show how a point cloud representation P leads to a smooth surface $S_f \in \mathcal{M}^{sph}$ using a spectral analysis of P .
- c) We define the Grasp Moduli Spaces \mathcal{G}^{sph} and \mathcal{G}^{rad} for m contacts on parametrized shapes in \mathcal{M}^{sph} and \mathcal{M}^{rad} respectively. We show how distances of spectral coefficients can be used to reason about nearby grasp/object configurations.
- d) We evaluate the success of a simple gradient ascent approach in \mathcal{G}^{rad} in order to locally optimize grasp quality on a fixed surface.
- e) We show how to efficiently transfer grasps in \mathcal{G}^{rad} between objects using a gradient based optimization and evaluate our approach in simulation.

The paper is structured as follows: In Sec. II, we discuss our motivation for this work and review common shape representations in robotic manipulation and the basics of spherical harmonics. In Sec. III, we discuss the theoretical aspects of our work, and in Sec. IV we present experimental results. We conclude in Sec. V and discuss future work.

II. BACKGROUND AND MOTIVATION

A. Grasp Synthesis

For a robot to grasp and manipulate objects, it is necessary to first determine an appropriate mathematical abstraction of this process. Grasp synthesis classically relies on a notion of grasp quality [4] and *force-closure* in particular (see [5] for a review). In order to determine if a grasp is force closed, one then typically assumes that all contacts between the robot hand and the object are point contacts, and that the geometry of the object is explicitly described – in practice usually by a triangle mesh. We hence think of a grasp g on a surface S with centre of mass z , with m contact points c_i and associated inward-pointing normal vectors n_i as a configuration

$$g = (c_1, \dots, c_m, n_1, \dots, n_m, z) \in \mathbb{R}^{3m} \times (\mathbb{S}^2)^m \times \mathbb{R}^3$$

and will use the L^1 grasp quality measure Q defined in [4] to determine if such a configuration is force closed given a surface friction coefficient $\mu > 0$. Q is defined to be the radius of the largest ball around the origin and entirely inside the grasp wrench space if such a sphere exists and zero otherwise. Q takes positive values if wrenches in arbitrary direction can be withstood by the grasp g and the grasp is considered more stable the larger Q is. State of the art approaches such as [6] then typically incorporate a grasp quality ranking with sampling and a decomposition of the object in order to return a list of viable grasp hypotheses.

B. Shape Representations for Grasping

Triangle meshes provide a commonly used object representation for robotic grasping. To determine contact configurations on a mesh, a common approach used *e.g.* by the popular simulation environment GraspIT [7] is based on sampling hand poses that establish contact with the object and to then rank the result by a grasp quality measure such as the previously mentioned Q function. The work of [8] also uses a sampling based approach to determine stable contact configurations. A reduced object representation based on the medial axes of a surface has been investigated in [9] for the purpose of grasp synthesis. An alternative approach is based on a hierarchical decomposition of objects into primitives such as spheres, boxes or quadrics. The work of [1] focusses on a box based decomposition in particular.

A problem with these representations is that it is not clear how a robot can utilize learned grasps in novel situations, *e.g.* when an object which is slightly different from previously investigated cases is encountered. Unlike problems in Computer Vision, where it is possible to obtain millions of training examples for a given learning problem, a robot will typically not grasp a given object more than a few dozen times and even simulated grasps are difficult to obtain in sufficient numbers so that the high-dimensional state-space consisting of robot hand configurations is densely sampled for each object instance.

A guiding motivation for our work is hence the desire to determine a representation allowing us to transfer grasps

between similar objects and to enable a statistical analysis of ‘similar’ grasps on ‘similar’ surfaces on a unified space of shapes and grasps. The work of [10] has taken a step in this direction, using point cloud matching and a combinatorial search to transfer grasps between objects. The approach of [11] instead decomposes an object into parts that can be studied separately. There, graspable parts are then matched between objects for the purpose of grasp transfer.

Note that, in both approaches, the surface geometry is represented using discrete surface descriptors. Our approach is instead based on a smooth object parametrization and extends the work of [2], where the space of surfaces with cylindrical coordinates, \mathcal{M}^{cyl} , was considered. The space \mathcal{M}^{cyl} is convex, allowing us to deform between any two surfaces S_1, S_2 using a curve $\gamma : [0, 1] \rightarrow \mathcal{M}^{cyl}$,

$$\gamma(t) = (1-t)S_1 + tS_2.$$

A further benefit of the convexity is the fact that any finite set of surfaces $X = \{S_1, \dots, S_n\}$ induces a finite dimensional shape subset $\text{Conv}(X)$ consisting of convex combinations of elements of X . In [2], this shape space is combined with a space of configurations of contact points to define a Grasp Moduli Space \mathcal{G}^{cyl} parametrizing both a shape and a set of contact points on the shape. A metric on \mathcal{G}^{cyl} can then be defined and one can continuously move between any two points in \mathcal{G}^{cyl} . A particular advantage of this approach is that there exist subsets of shapes on which grasp synthesis can be carried out very efficiently, such as the surfaces of revolution $\mathcal{M}^{rev} \subset \mathcal{M}^{cyl}$ considered in [2]. Using a deformation in \mathcal{G}^{cyl} , these grasps can then be generalised to nearby surfaces which are not themselves rotationally symmetric.

Noise can also naturally be described in \mathcal{G}^{cyl} since we are able to describe ‘shape-noise’ and ‘contact point configuration noise’. Three important cases are 1) when the shape is precisely perceived, but the hand cannot be controlled sufficiently precisely, 2) when a slightly deformed shape is encountered but the grasp was planned on the original internal model of the shape, and 3) when both the shape and the final hand position are not precisely the same as an internal representation had predicted.

We will utilize spherical harmonics to obtain a smooth surface representation which have surfaced for manipulation in the work of [12], where the resulting smooth parametrization was used to control a rolling interaction between a simple gripper and an object. Based on the resulting surface parametrizations, we discuss a shape space which generalises \mathcal{M}^{cyl} since, in principle, any closed surface of genus zero can be modelled given dense samples of points from the object’s surface.

C. Spherical Harmonics

In order to extract smooth surfaces from point cloud data, we shall make use of a spectral analysis on the sphere which provides an analogue of classical Fourier analysis. Consider the space $L^2(\mathbb{S}^2)$ of integrable functions on the sphere $\mathbb{S}^2 \subset \mathbb{R}^3$. We choose a coordinate chart $p(\theta, \varphi) = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$ on the sphere, with $(\theta, \varphi) \in [0, \pi] \times [0, 2\pi)$. For integers l, m , the real valued spherical

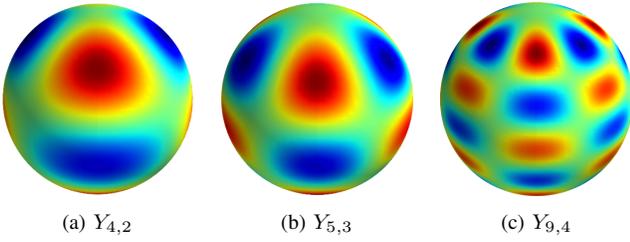


Fig. 1: Illustration of some eigen-functions of Δ on \mathbb{S}^2 .

harmonic function of degree l and order m , $|m| \leq l$, is defined by

$$Y_{l,m}(\theta, \varphi) = \begin{cases} c_{l,m} P_l^{|m|}(\cos \theta) \sin(|m|\varphi) & -l \leq m \leq -1 \\ \frac{c_{l,m}}{\sqrt{2}} P_l^0(\cos \theta) & m = 0 \\ c_{l,m} P_l^m(\cos \theta) \cos(m\varphi) & 1 \leq m \leq l, \end{cases}$$

where $c_{l,m} = \sqrt{\frac{2l+1}{2\pi} \frac{(l-|m|)!}{(l+|m|)!}}$ and P_l^m denotes the associated Legendre polynomial of order m and degree l . The functions $Y_{l,m}$ arise as eigen-functions of the Laplacian $\Delta = \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} (\sin \theta \frac{\partial}{\partial \theta}) + \frac{1}{\sin^2 \theta} \frac{\partial^2}{\partial \varphi^2}$ on \mathbb{S}^2 [3]. For fixed l , there are $2l + 1$ eigenfunctions $\{Y_{l,m} : m \in \{-l, \dots, l\}\}$ satisfying

$$\Delta Y_{l,m} = \lambda_l Y_{l,m},$$

where $\lambda_l = l(l+1)$ denotes the corresponding eigenvalue. These eigen-functions, for all l, m , form an infinite orthonormal basis for $L^2(\mathbb{S}^2)$ with respect to the standard inner product $\langle f, g \rangle = \int_{\mathbb{S}^2} fg \, d\text{Vol}$, for $f, g \in L^2(\mathbb{S}^2)$ and where $d\text{Vol} = \sin \theta d\theta d\varphi$. Fig. 1 displays a few of these eigen-functions. Any $f \in L^2(\mathbb{S}^2)$ can hence be expanded as

$$f(\theta, \varphi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{l,m} Y_{l,m}(\theta, \varphi),$$

for some $a_{l,m} \in \mathbb{R}$. Suppose now that we have samples $p_i = p(\theta_i, \varphi_i) \in \mathbb{S}^2$ on the sphere and associated noisy function values $f_i \in \mathbb{R}$, for $1 \leq i \leq n$, of some underlying function $f : \mathbb{S}^2 \rightarrow \mathbb{R}$. We can approximately reconstruct f using a function $h \in L^2(\mathbb{S}^2)$ such that h lies in the direct sum H_L of the eigen-spaces of Δ up to degree L . The vector space H_L has dimension $(L+1)^2$ and any such h is of the form

$$h(\theta, \varphi) = \sum_{l=0}^L \sum_{m=-l}^l \beta_{l,m} Y_{l,m}(\theta, \varphi),$$

for some $\beta_{l,m} \in \mathbb{R}$. A least-squares fit, minimizing the error term

$$\sum_{i=1}^n |h(\theta_i, \varphi_i) - f_i|^2$$

is given as follows (see [13] for details): let $\mathbf{f} = (f_1, \dots, f_n)$, $\beta_l = (\beta_{l,-l}, \dots, \beta_{l,l})$,

$$\mathbf{Y}_l = \begin{bmatrix} Y_{l,-l}(\theta_1, \varphi_1) & \dots & Y_{l,l}(\theta_1, \varphi_1) \\ \vdots & \ddots & \vdots \\ Y_{l,-l}(\theta_n, \varphi_n) & \dots & Y_{l,l}(\theta_n, \varphi_n) \end{bmatrix}.$$

Furthermore, define $\mathbf{Y} = [\mathbf{Y}_0, \dots, \mathbf{Y}_L]$ and $\beta = (\beta_0^t, \dots, \beta_L^t)^t$. Then, as explained in more detail in [13], a choice of coefficients given by $\beta = (\mathbf{Y}^t \mathbf{Y})^{-1} \mathbf{Y}^t \mathbf{f}$ yields a solution $h \in H_L$ minimizing the error term. This approach is part of the theory of spline based regression which can be extended by a regularization term [14]. The above procedure hence yields a regression methodology for functions defined on spheres. As can be seen in examples in [3], the first $L = 70$ eigen-spaces are usually sufficient to capture even minute details of f . The work of [15] uses the above approach to find a spherical parametrization for surface meshes. As a first step, the mesh-vertex x, y, z coordinates are parametrized as functions on \mathbb{S}^2 using the methods outlined in [16]. The resulting functions $x, y, z : \mathbb{S}^2 \rightarrow \mathbb{R}$ are then approximated using an expansion in terms of $Y_{l,m}$ as above. Given a sufficiently densely sampled general point cloud of a closed surface with genus zero (*i.e.* without holes), one can also determine an embedding mapping the point cloud onto a sphere. One approach for this is based on energy minimization [17], while a recent method uses an approximated Laplace-Beltrami operator for this purpose [18], see [19] for an overview. Once the x, y, z coordinates of the sampled points (x_i, y_i, z_i) are mapped to corresponding θ_i, φ_i values, the above regression method can be applied and yields a smooth parametrization $(x, y, z) : \mathbb{S}^2 \rightarrow \mathbb{R}^3$ of the point cloud. In neuroanatomy, [13] recently used a spherical harmonics expansion of highly curved brain surfaces to characterize cortical thickness. There, a subject's brain surface is first mapped onto a sphere and regression is then performed on this transformed data based on the least squares approach outlined above.

III. THEORETICAL FRAMEWORK

We now describe how to extract a suitable object representation from point cloud data. We then formalize and study the shape space which we will use in this paper.

A. Spaces of Shapes and Spherical Coordinates

In this work, we are interested in smooth closed embedded surfaces in \mathbb{R}^3 of genus zero. We make the following definition:

Definition 3.1: Given $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$, we denote the image by $S_f = \{f(p) : p \in \mathbb{S}^2\} \subset \mathbb{R}^3$. We define

$$\mathcal{M}^{sph} = \{f \in \mathcal{C}^\infty(\mathbb{S}^2, \mathbb{R}^3) : f \text{ is an smooth embedding.}\}$$

We furthermore denote by \mathcal{M}^{rad} those $f \in \mathcal{M}^{sph}$ for which there exists a radial function $r : \mathbb{S}^2 \rightarrow \mathbb{R}_{>0}$ such that

$$f(p(\theta, \varphi)) = r(\theta, \varphi) (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta),$$

for all $(\theta, \varphi) \in [0, \pi] \times [0, 2\pi)$.

The set \mathcal{M}^{sph} describes parametrizations of closed embedded smooth surfaces of genus zero, while $\mathcal{M}^{rad} \subset \mathcal{M}^{sph}$ contains parametrizations of simpler such surfaces which can be described by a single radial coordinate function $r : \mathbb{S}^2 \rightarrow \mathbb{R}_{>0}$. The surfaces defined by such radial functions yield an interesting subset $\mathcal{M}^{rad} \subset \mathcal{M}^{sph}$ which we shall explore in this work in particular. Recall that the image S_f

of a smooth embedding does not have self-intersections and yields a smooth closed surface suitable for representing a variety of real world objects. We can mirror the observations made about the space of surfaces with cylindrical coordinates in [2] also in the case of \mathcal{M}^{rad} :

Lemma 3.2: For any $f_1, f_2 \in \mathcal{M}^{rad}$ and $\alpha_1, \alpha_2 > 0$, we have $\alpha_1 f_1 + \alpha_2 f_2 \in \mathcal{M}^{rad}$. In particular, \mathcal{M}^{rad} is convex.

Proof: Let $f_1, f_2 \in \mathcal{M}^{rad}$ with corresponding radial functions $r_1, r_2 : \mathbb{S}^2 \rightarrow \mathbb{R}_{>0}$. Then, for $\alpha_1, \alpha_2 > 0$, $f = \alpha_1 f_1 + \alpha_2 f_2$ has radial function $\alpha_1 r_1 + \alpha_2 r_2$ which is smooth and positive. By a simple explicit computation of f and Jf in smooth coordinate charts covering \mathbb{S}^2 , one can verify that f is injective and that Jf has full rank everywhere. Hence f is an injective immersion. Since \mathbb{S}^2 is compact, it follows that f is in fact a smooth embedding. Hence $f \in \mathcal{M}^{rad}$. The statement about convexity follows at once. ■

Note that the above does *not* state that the shapes in \mathcal{M}^{rad} themselves have to be convex, but that the shape space \mathcal{M}^{rad} is. The above simple observation implies in particular that, given any finite set of smooth functions $X = \{f_1, \dots, f_n\} \subset \mathcal{M}^{rad}$, $\text{Conv}(X) \subset \mathcal{M}^{rad}$ and any element of $\text{Conv}(X)$ yields a valid smooth embedded surface. $\text{Conv}(X)$ hence yields a natural finite dimensional subspace of surfaces induced by a finite set of example surfaces. Similarly, we can consider $\text{Cone}(X) = \{\alpha f : \alpha > 0, f \in \text{Conv}(X)\}$.

We can deform between any two surfaces $f_1, f_2 \in \mathcal{M}^{rad}$ using $f(t, p) = (1-t)f_1(p) + tf_2(p)$, $t \in [0, 1]$, $p \in \mathbb{S}^2$ and, more generally, interpolate between any finite set of example surfaces. For the larger set \mathcal{M}^{sph} , one can still consider such convex combinations, but care has to be taken to verify that the resulting surfaces are still elements of \mathcal{M}^{sph} to make sure no degeneracies or self-intersections occur.

Distances between Shape Parametrizations: Note that, for $f_1, f_2 \in \mathcal{M}^{sph}$, we can define a distance function $d(f_1, f_2) = (\int_{\mathbb{S}^2} |f_1 - f_2|^2 d\text{Vol})^{\frac{1}{2}}$, where $d\text{Vol} = \sin \theta d\theta d\varphi$. If $f_1, f_2 \in \mathcal{M}^{rad}$ with corresponding radial functions r_1, r_2 , then in fact

$$d(f_1, f_2)^2 = \int_{\mathbb{S}^2} |f_1 - f_2|^2 d\text{Vol} = \int_{\mathbb{S}^2} |r_1 - r_2|^2 d\text{Vol}.$$

The above formula allows us to quantify the similarity between the corresponding parametrized surfaces S_{f_1}, S_{f_2} and we shall use the notation $d(f_1, f_2), d(r_1, r_2)$ interchangeably.

B. Spherical Harmonics and Point Clouds

Let us here describe our methodology for obtaining a surface parametrization $f \in \mathcal{M}^{sph}$ or $f \in \mathcal{M}^{rad}$ for a given input point cloud $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^3$ sampled from some genus zero surface $S \subset \mathbb{R}^3$. Such point clouds arise for example from laser range scanners or from a structured light sensor such as the kinect sensor. As mentioned previously, there are several methods available to assign spherical coordinates for each point in P [17], [18], [19] and these methods apply even for complicated non-convex objects. Once corresponding (θ_i, φ_i) coordinates are found, each of the functions $x, y, z : \mathbb{S}^2 \rightarrow \mathbb{R}$ can then be recovered using the presented methods. In the case where P was sampled from a surface

S_f with $f \in \mathcal{M}^{rad}$, corresponding θ_i, φ_i coordinates can be determined by computing the standard spherical coordinates of each $p \in P$ directly. Furthermore, instead of finding a regression solution for the x, y, z coordinates separately, one can then directly compute the standard radial coordinate r_i for each $p_i \in P$ and determine a radial function $r : \mathbb{S}^2 \rightarrow \mathbb{R}$ approximating the data $(r_i, \theta_i, \varphi_i)$ using eigen-spaces up to degree $L > 0$. Since the objects considered for grasping in this paper can be modelled as elements of \mathcal{M}^{rad} , we will employ this simplified approach here and defer a full study of more general objects to later work. The function $r : \mathbb{S}^2 \rightarrow \mathbb{R}$ is then given by

$$r(\theta, \varphi) = \sum_{l=0}^L \sum_{m=-l}^l \beta_{l,m} Y_{l,m}(\theta, \varphi),$$

and we denote by $\beta(r) = (\beta_{0,0}, \beta_{1,-1}, \dots, \beta_{L,L})$ the vector of coefficients characterizing r . Let r_1, r_2 be two such functions with coefficients $\beta(r_1), \beta(r_2)$. Observe that, since $\langle Y_{l,m}, Y_{l',m'} \rangle = \delta_{l,l'} \delta_{m,m'}$, where δ denotes the Kronecker delta function, we can quickly compute the distance d between the resulting surfaces using the vector norm:

$$d(r_1, r_2)^2 = \int_{\mathbb{S}^2} |r_1 - r_2|^2 d\text{Vol} = \|\beta(r_1) - \beta(r_2)\|^2.$$

Fig. 2 displays a point cloud P and a resulting reconstructed smooth surface parametrization $f \in \mathcal{M}^{rad}$ which is obtained by using all eigen-spaces of Δ up to degree L . As L is increased, more and more detail about the surface is captured. Note that – unlike common mesh representations for objects – the reconstructed surface parametrization f is smooth and derivatives of arbitrary order can be computed.

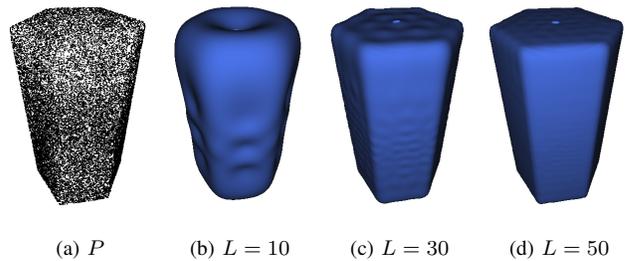


Fig. 2: Point cloud P with 25000 points and reconstructed surfaces, where all eigen-spaces up to degree L are used.

C. Grasp Moduli Spaces and Spherical Coordinates

A long term goal of our work is the construction of a shape and grasp configuration space which allows us to continuously move between shape/grasp configurations and where we can employ a) natural optimization methods for smooth objective functions and b) probabilistic methods which incorporate shape and grasp uncertainties. We now take a step in this direction by formalizing a space of grasps on \mathcal{M}^{sph} and a natural grasp/shape distance function.

Definition 3.3: The *Point-Contact Grasp Moduli Space* for m point-contacts on \mathcal{M}^{sph} is given by $\mathcal{G}^{sph} = \mathcal{M}^{sph} \times (\mathbb{S}^2)^m$. The corresponding space for \mathcal{M}^{rad} is given by $\mathcal{G}^{rad} = \mathcal{M}^{rad} \times (\mathbb{S}^2)^m$.

A point $(f, (p_1, \dots, p_m)) \in \mathcal{G}^{sph}$ now parametrizes both a surface $S = S_f$ and a grasp $g(f, (p_1, \dots, p_m)) = (c_1, \dots, c_m, n_1, \dots, n_m, z)$, where $c_i = f(p_i)$ denotes the i^{th} contact point, n_i the surface normal of S_f at c_i and z the corresponding centre of mass of S_f . As in [2], we can define a metric on grasp contact configurations as follows: consider $g, g' \in \mathbb{R}^{3m} \times (\mathbb{S}^2)^m \times \mathbb{R}^3$,

$$g = (c_1, \dots, c_m, n_1, \dots, n_m, z)$$

$$g' = (c'_1, \dots, c'_m, n'_1, \dots, n'_m, z')$$

and define

$$d(g, g') = \max_i (d_{\mathbb{R}^3}(c_i, c'_i), d_{\mathbb{R}^3}(n_i, n'_i), d_{\mathbb{R}^3}(z, z')).$$

We can in turn combine this with our metric on \mathcal{M}^{sph} to define a metric on \mathcal{G}^{sph} by $d((f, (p_1, \dots, p_m)), (f', (p'_1, \dots, p'_m))) = wd(f, f') + (1 - w)d(g, g')$, for some fixed $w \in [0, 1]$. Here, $g = g(f, (p_1, \dots, p_m))$ and $g' = g(f', (p'_1, \dots, p'_m))$. Using this metric, we can quantify how far a given grasp/surface configuration is from a different such configuration. The w parameter can be chosen to weigh the relative importance of these distances. Note also that we can continuously deform $(f, (p_1, \dots, p_m)) \in \mathcal{G}^{rad}$ to $(f', (p'_1, \dots, p'_m)) \in \mathcal{G}^{rad}$, by a curve $\gamma = (u, v) : [0, 1] \rightarrow \mathcal{G}^{rad}$, where e.g. $u(t) = (1 - t)f + tf'$ and v deforms each $p_i \in \mathbb{S}^2$ along a spherical geodesic towards p'_i . The same approach can be attempted in \mathcal{G}^{sph} , but one has to take care that $S_{u(t)} \in \mathcal{M}^{sph}$ for all $t \in [0, 1]$.

IV. EXPERIMENTS

We now describe experiments which utilize the previously discussed framework in various ways in order to generate force closed grasp hypotheses and to explore the resulting moduli space. Details about our approach and our implementation of spherical harmonics based regression can be found on the first author's website <http://www.csc.kth.se/~fpokorny>.

A. Object Data and Shape Space

We work with the collection of real world objects depicted in Fig. 3. For each object, we generate a point cloud P_j of 25000 uniformly sampled points on the object's surface which forms the basis of our approach. For this purpose, we used the laser scan data of [20] for all surfaces but S_2, S_7, S_{11}, S_{12} in Fig. 4, for which we generated uniform samples from our own mesh models.

This data is then centred at the mean of the point cloud which we assume to be the centre of mass of the object, and the spherical coordinates for each data-point are computed. We utilize the spherical regression technique described in Sec. III-B and all eigen-spaces up to degree $L = 50$ to obtain a respective smooth radial function $r : \mathbb{S}^2 \rightarrow \mathbb{R}_{>0}$ for each point cloud which then results in a surface S_f , $f(\theta, \varphi) = r(\theta, \varphi)(\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta)$. Fig. 4 displays the resulting surfaces S_1 up to S_{12} . A more detailed view of the point cloud for S_8 can also be seen in Fig. 2.



Fig. 3: A picture of the objects used in our experiments.

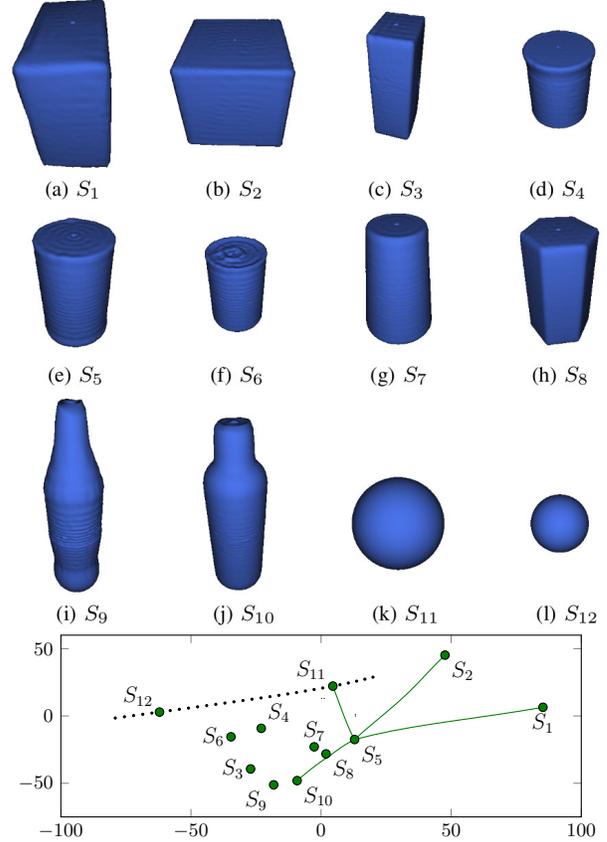


Fig. 4: The first 3 rows display smooth parametrized surfaces determined from point clouds (25000 points each) of our objects from Fig. 3. The last row displays the MDS projection of β vectors parametrizing these objects in \mathcal{M}^{rad} (green dots), a sequence of spheres (black dots) and example deformations (solid lines).

Each surface corresponds to a β coefficient vector in \mathbb{R}^{2601} . In the bottom part of Fig. 4, we visualize the distribution of the β vectors for our 12 surfaces, indicated by green dots. To enrich the set of available shapes, we also created 25000 uniformly distributed random points and respective smooth reconstructions from 30 spheres with radii between 3cm and 6cm and we added examples of shape deformations from shape S_5 to S_1, S_2, S_{10}, S_{11} , where a deformation between S_i and S_j is obtained by a convex combination of the respective vectors β_i, β_j : $\beta(t) = (1 - t)\beta_i + t\beta_j$. On the resulting set of β vectors, we then performed manifold multidimensional scaling (MDS) to obtain Fig. 4. Note that,

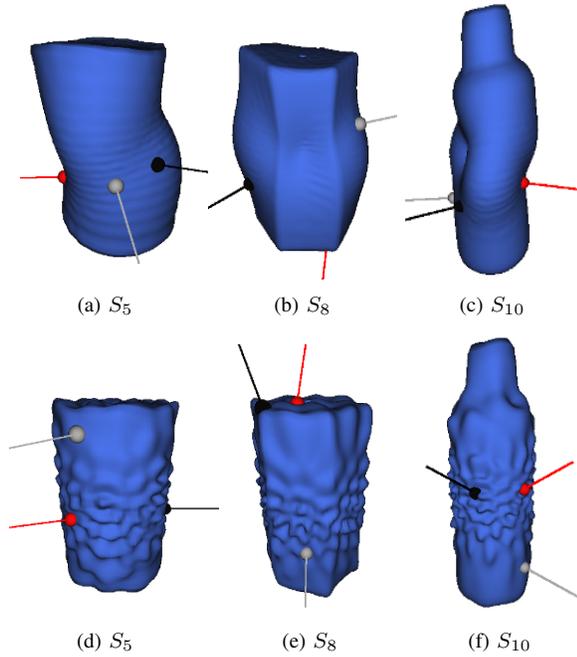


Fig. 5: Random samples of points in \mathcal{G}^{rad} , where Gaussian noise is added to β_0, \dots, β_9 in the top row and to $\beta_{499}, \dots, \beta_{599}$ in the bottom row.

as S_5 is deformed, it passes the neighbourhood of the shapes S_7 and S_8 . The deformations correspond to straight lines in \mathbb{R}^{2601} . The sequence of spheres is furthermore also clearly visible as a sequence of nearby points passing S_{11} and S_{12} .

B. Compatibility with probabilistic approaches

Note that the shape-space component of \mathcal{G}^{sph} as approximated by the first L eigen-spaces of Δ is simply corresponding to a $(L+1)^2$ dimensional vector space containing β coefficients. It is hence clear that the full vector-space based machinery of machine learning and statistical analysis is available on this space. Additionally, appropriate distributions on spheres exist, endowing the coordinate part of \mathcal{G}^{sph} with natural probability distributions. One can hence consider building *e.g.* mixture models on our moduli space. In future work, we intend to explore this direction. One could for example consider modelling task constraints in grasping probabilistically using such distributions. In order to illustrate sampling in \mathcal{G}^{sph} , we created the deformed object/grasp configurations in Fig. 5 by sampling three random (θ, φ) grasp contact coordinates and by adding Gaussian noise to the first 10 dimensions of β in the top row and to dimensions 499 to 599 in the bottom row. As we can observe, in a situation mirroring the familiar case of Fourier analysis on the circle, modifications of the first few β coefficients correspond to large scale object properties, while the higher dimensions correspond to higher ‘frequencies’.

C. Initial grasp synthesis

We work with a Schunk dexterous hand (SDH, 7 DOF) displayed in Fig. 6 for which we use a simplified collision model displayed in the top left part of that figure. To detect collisions with our reconstructed smooth surfaces, we used a

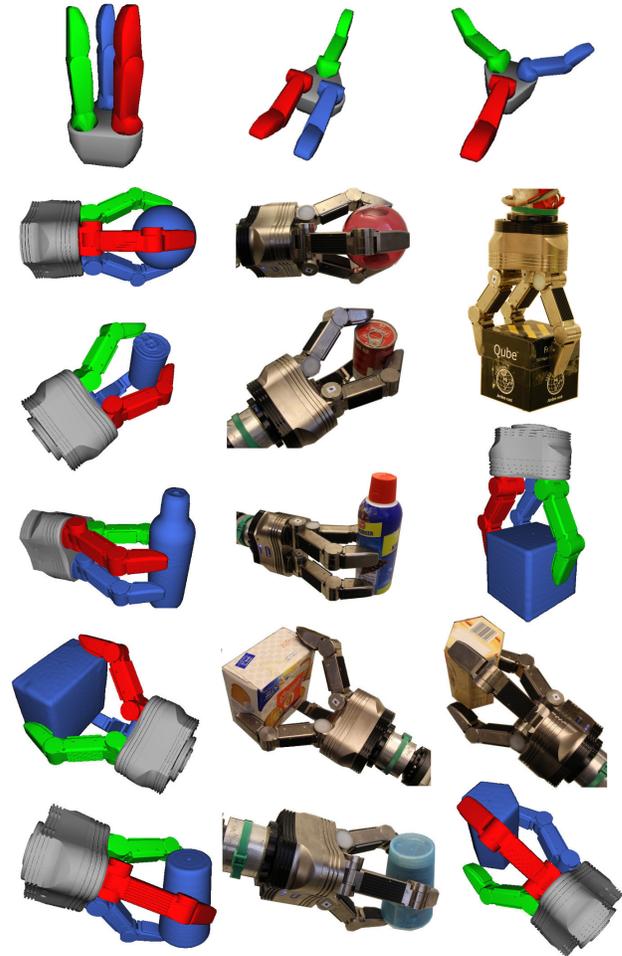


Fig. 6: Top row: SDH hand collision model and two pre-shapes. Remaining rows: stable grasps generated in simulation and their execution in reality (background whitened).

regular mesh on $(\theta, \varphi) \in [0, \pi] \times [0, 2\pi]$ with 10000 vertices and mapped it onto the surface using $f(\theta, \varphi)$. These meshes are also used in Fig.4 for visualization. We sampled SDH hand positions oriented towards the origin and on the upper hemisphere of varying radii in an approach similar to the one used by GraspIT [7]. The hand was put either in a parallel pre-shape with open fingers or in a pre-shape where the fingers are equally spread out (see the two figures to the right in the first row of Fig. 6). We then closed the fingers until the finger-tips were less than 1mm away from the object. Configurations where the fingers did not reach the surface vicinity, or where other hand-parts were less than 5mm away from the object or closer than 5mm to each other were discarded. The closest points p_1, p_2, p_3 on the object and from the finger-tips were calculated and corresponding unit normals n_1, n_2, n_3 of the surface S_f were computed based on our smooth representation by normalizing $\frac{df}{d\theta} \times \frac{df}{d\varphi}$. From an initial set of 1400 random hand-positions per surface, we obtained a set of more than 200 stable grasps per object which we ranked by grasp quality. Let us denote by G_i the 50 best such grasps and by \hat{G}_i the best 10 grasps on surface S_i . To verify that grasps which were stable in simulation also

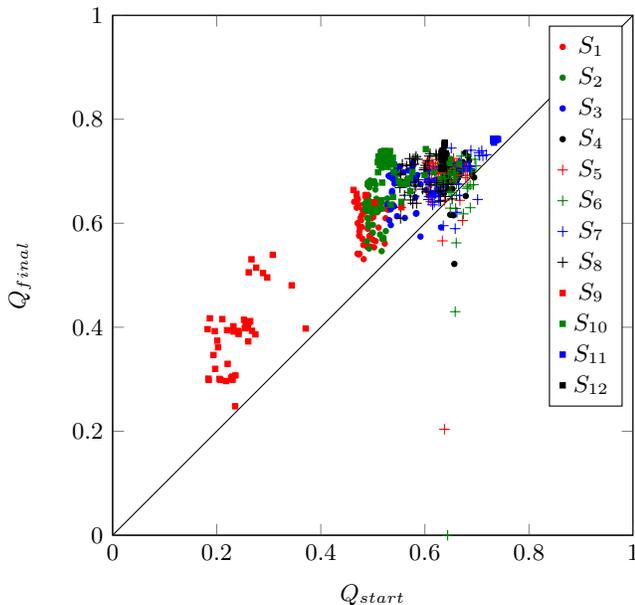


Fig. 7: Initial vs. final grasp quality with gradient ascent

corresponded to stable grasps in reality, we performed a short evaluation where two simulated stable grasps on each surface were executed on a real Schunk (SDH) hand attached to a Kuka arm. 21 of the resulting grasps were stable, while 1 failure occurred on each of the objects S_1, S_3, S_6 – most likely due to errors in the positioning of the hand. The supplementary video contains more details on this evaluation and examples from these experiments are also displayed in Fig. 6 alongside the corresponding simulated grasps.

D. Point-contact based grasp synthesis and transfer on \mathcal{G}^{rad}

The synthesis and optimization of stable point-contact configurations are important sub-problems in grasp synthesis since stable configurations can serve as goals for a planning algorithm positioning the robot hand.

We now study the problem of grasp optimization given a stable starting grasp. For this purpose, we consider a grasp g on a surface S_f with three contact points. Each such grasp is determined by f and $(\theta, \varphi) \in [0, \pi]^3 \times [0, 2\pi]^3 \subset \mathbb{R}^6$. The grasp quality Q can hence be thought of as a function of θ, φ which we attempt to maximize. We propose a straightforward gradient ascent approach as follows: we approximate $(\nabla Q(x))_i \approx \frac{1}{\delta}(Q(x + \delta e_i) - Q(x))$, where $e_i \in \mathbb{R}^6$ denotes the i -th basis vector, $i \in \{1, \dots, 6\}$, with $\delta = 10^{-5}$ as in [2]. Starting from an initial grasp $x_0 \in [0, \pi]^3 \times [0, 2\pi]^3$, we update $x_{n+1} = x_n + 0.01 \text{clamp}(\nabla Q(x_n))$, where $\text{clamp}(x) = x$ if $\|x\| \leq 1$ and $\frac{x}{\|x\|}$ otherwise. We applied 1000 steps of this approach to each grasp in G_i for each surface S_i . The resulting initial and final grasp qualities are summarized in Fig. 7 with grasps on different surfaces indicated by differing symbols and colours. Note that grasp quality is typically improved significantly, but that, in some cases, our simple fixed step-length gradient ascent does not converge and can in fact result in final grasps which have

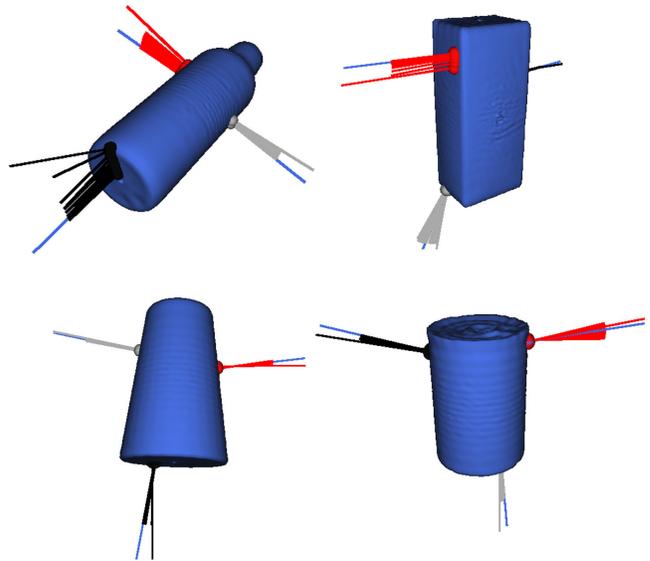


Fig. 8: Initial grasp (long black, grey and red normals) and changes during the gradient ascent (smaller lines) are displayed. The final contact configurations are depicted as long blue lines. The grasp quality improved in all but the bottom right example.

worse quality than the initial grasp. This behaviour could be partially avoided by using an adaptive step size gradient ascent approach instead. Fig. 8 displays successful gradient ascent results with an improvement in Q from 0.51, 0.54, 0.62 and 0.62 to 0.72 (top left), 0.64 (top right), 0.70 (bottom left) and 0.67 (bottom right) respectively. Note how a large adjustment to the grasp configuration is made in the top left plot while in the bottom right example only a small change occurs during our gradient ascent. We can think of this gradient ascent as forming a sequence of points in \mathcal{G}^{rad} , where the shape part is held fixed. Let us remark furthermore what we could use our approach with any grasp grasp quality function for which we can approximate gradients, including quality functions taking into account additional constraints.

E. Optimization-based grasp transfer in \mathcal{G}^{rad}

Next, we investigate the ability of our approach to be used to transfer stable grasps from one object to another one. We consider, for $i, j \in \{1, \dots, 12\}$, $i \neq j$ the curve $\gamma : [0, 1] \rightarrow \mathcal{M}^{rad}$ obtained by scaling the parametrizations f_i, f_j corresponding to the surfaces S_i and S_j by $\gamma(t) = (1-t)f_i + tf_j$. This deformation in turn corresponds to linear interpolation between the respective β -coefficients of S_i, S_j . The naïve approach to transfer a grasp to a new surface is obtained by simply fixing the θ, φ coordinates of the contact points as the underlying surface deforms. We study an alternative approach, as introduced in [2] for \mathcal{G}^{cyl} , and start with a grasp $g \in \hat{G}_i$ on S_i , apply 100 steps of our gradient ascent on $S_{f_i} = S_{\gamma(0)}$, then we move the grasp to the surface $S_{\gamma(0.1)}$ and transfer the grasp to the slightly deformed surface by keeping the same θ, φ coordinates per contact. We then apply again 100 steps of our gradient ascent and move to $S_{\gamma(0.2)}$ etc. until we reach a grasp on $S_{\gamma(1)} = S_{f_j}$ which is again optimized with 100 steps of our

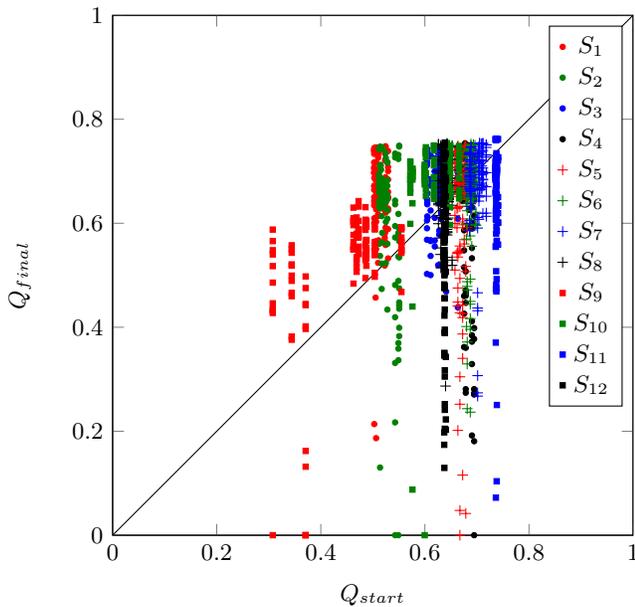


Fig. 9: We display initial vs. final grasp quality when transferring grasps from one surface to another.

gradient ascent. We applied this procedure for every one of the grasps in \hat{G}_i and every starting surface S_i and target surface $S_j \neq S_i$, resulting in 110 grasps per starting surface and 1320 grasp transfers in total. Note that we concentrate on the problem of grasp transfer from a point-contact point of view here. This approach can then be integrated at a later stage with a motion planning approach to adapt the hand as the point contact configuration is deformed. The results of our deformation based transfer are summarized in Fig. 9. As one might expect, the grasp quality on the final target surface is not always positive since our simple optimization method can get stuck in local maxima for this non-convex optimization problem. However, the vast majority of the final grasps is indeed stable and of good grasp quality after the grasp is transferred. Among the 1320 transferred grasps only 9 had a grasp quality of less than 0.01 which we consider unstable. Of 110 transfers per starting surface only 2, 1, 1, 4 and 1 such unstable grasps occurred for the starting surface S_2 , S_4 , S_5 , S_9 and S_{10} respectively, while all 110 grasp transfers succeeded for each of the remaining starting surfaces. The figure on the first page illustrates two deformations in \mathcal{G}^{rad} during this optimization. The contact normals are indicated by red, grey and green lines.

V. CONCLUSION AND OPEN PROBLEMS

In this work, we have further developed our long term aim of formulating an object/grasp representation which allows for a deformation based grasp synthesis in a combined space of grasps and shapes. We have shown that commonly available point cloud data can be used with our approach and that it is naturally compatible with gradient based optimization methods for point contact synthesis. Additionally, our approach is compatible with probabilistic modelling techniques since probability distributions can naturally be defined on our Grasp Moduli Space \mathcal{G}^{sph} . Fundamental challenges however

remain. It is clear that our smooth parametrization will deteriorate with partial point cloud data. An investigation of the connection of our approach to Gaussian Process Regression [14] and a suitable probability distribution on our shape space to accommodate missing data while avoiding unrealistic shape reconstructions are potential future research directions. Similarly, our method currently requires an initial matching of the object's rotation, e.g. using an iterative closest point method. It would be interesting to further study rotationally invariant methods which are compatible with our deformation based framework. Finally, it is clear that our gradient ascent approach is not guaranteed to converge to a global optimum. We hence intend to also explore alternative numerical approaches for optimizing grasp quality within our deformation based Grasp Moduli Space framework.

REFERENCES

- [1] K. Huebner, S. Ruthotto, and D. Kragic, "Minimum volume bounding box decomposition for shape approximation in robot grasping," in *IEEE ICRA*, 2008, pp. 1628–1633.
- [2] F. T. Pokorny, K. Hang, and D. Kragic, "Grasp moduli spaces," in *Proceedings of Robotics: Science and Systems*, 2013.
- [3] M. K. Chung, *Computational Neuroanatomy: The Methods*. World Scientific, 2013.
- [4] C. Ferrari and J. Canny, "Planning optimal grasps," *IEEE ICRA*, pp. 2290–2295, 1992.
- [5] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *IEEE ICRA*, 2000, pp. 348–353 vol.1.
- [6] J.-P. Saut and D. Sidobre, "Efficient models for grasp planning with a multi-fingered hand," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 347–357, 2012.
- [7] A. T. Miller and P. K. Allen, "Grasplit: A versatile simulator for robotic grasping," *IEEE Robotics and Automation Magazine*, vol. 11, pp. 110–122, 2004.
- [8] C. Borst, M. Fischer, and G. Hirzinger, "Grasping the dice by dicing the grasp," in *IEEE/RSJ IROS*, 2003, pp. 3692–3697 vol.4.
- [9] M. Przybylski, T. Asfour, and R. Dillmann, "Planning grasps for robotic hands using a novel object representation based on the medial axis transform," in *IEEE/RSJ IROS*, 2011, pp. 1781–1788.
- [10] U. Hillenbrand and M. A. Roa, "Transferring functional grasps through contact warping and local replanning," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 2963–2970.
- [11] R. Detry, C. H. Ek, M. Madry, and D. Kragic, "Learning a dictionary of prototypical grasp-predicting parts from grasping experience," in *IEEE International Conference on Robotics and Automation*, 2013.
- [12] A. Bicchi and A. Marigo, "Dexterous grippers: Putting nonholonomy to work for fine manipulation," *The International Journal of Robotics Research*, vol. 21, no. 5-6, pp. 427–442, 2002.
- [13] M. K. Chung, K. M. Dalton, L. Shen, A. C. Evans, and R. J. Davidson, "Weighted fourier series representation and its application to quantifying the amount of gray matter," *Medical Imaging, IEEE Transactions on*, vol. 26, no. 4, pp. 566–581, 2007.
- [14] G. Wahba, *Spline models for observational data*. Siam, 1990, vol. 59.
- [15] L. Shen and M. K. Chung, "Large-scale modeling of parametric surfaces using spherical harmonics," in *Third Int. Symp. on 3D Data Processing, Visualization, and Transmission*. IEEE, 2006.
- [16] M. S. Floater and K. Hormann, "Surface parameterization: a tutorial and survey," in *Advances in multiresolution for geometric modelling*. Springer, 2005, pp. 157–186.
- [17] M. Alexa, "Recent advances in mesh morphing," in *Computer graphics forum*, vol. 21, no. 2. Wiley Online Library, 2002, pp. 173–198.
- [18] M. Zwicker and C. Gotsman, "Meshing point clouds using spherical parameterization," in *Proc. of the First Eurographics conf. on Point-Based Graphics*. Eurographics Association, 2004.
- [19] C. Gotsman, X. Gu, and A. Sheffer, "Fundamentals of spherical parameterization for 3d meshes," in *ACM Transactions on Graphics (TOG)*, vol. 22, no. 3. ACM, 2003, pp. 358–363.
- [20] R. Dillmann, "Kit objectmodels web database," 2013. [Online]. Available: <http://i61p109.ira.uka.de/ObjectModelsWebUI/>