

Grasp Moduli Spaces

Florian T. Pokorny, Kaiyu Hang and Danica Kragic

Centre for Autonomous Systems/CVAP, KTH Royal Institute of Technology

{fpokorny, kaiyuh, danik}@csc.kth.se

Abstract—We present a new approach for modelling grasping using an integrated space of grasps and shapes. In particular, we introduce an infinite dimensional space, the Grasp Moduli Space, which represents shapes and grasps in a continuous manner. We define a metric on this space allowing us to formalize ‘nearby’ grasp/shape configurations and we discuss continuous deformations of such configurations. We work in particular with surfaces with cylindrical coordinates and analyse the stability of a popular L^1 grasp quality measure Q_l under continuous deformations of shapes and grasps. We experimentally determine bounds on the maximal change of Q_l in a small neighbourhood around stable grasps with grasp quality above a threshold. In the case of surfaces of revolution, we determine stable grasps which correspond to grasps used by humans and develop an efficient algorithm for generating those grasps in the case of three contact points. We show that sufficiently stable grasps stay stable under small deformations. For larger deformations, we develop a gradient-based method that can transfer stable grasps between different surfaces. Additionally, we show in experiments that our gradient method can be used to find stable grasps on arbitrary surfaces with cylindrical coordinates by deforming such surfaces towards a corresponding ‘canonical’ surface of revolution.

I. INTRODUCTION

A lot of attention in the robotics community has been focussed on finding good representations that enable a robot to grasp novel objects. Since the problem of determining physically stable grasps is such a challenging one, it has been popular to think of the object and grasp representation separately in order to gain further insights. Object representations based on box-models [10], shape primitives [13], the medial axis [17] and string based structure representations [12] have been successfully explored. A popular approach is to first develop a grasp heuristic resulting in a set of grasp hypotheses on a particular object and to then rank the synthesized grasps by a grasp quality measure [4, 5, 7] depending on the contacts between the object and the robot hand. Most existing methods then proceed by evaluating their proposed grasp hypotheses framework on a discrete set of objects [10, 13].

In this work, we do not think of a collection shapes in a discrete manner, but instead take a step towards understanding grasps on families of shapes varying *continuously* with some parameter. Many real world objects can be thought of as varying continuously: besides simple properties such as height, width etc., we can imagine that every-day object classes such as bottles, tools, cars, *etc.* depend continuously on some ‘style’ parameters. In order to state a mathematically concise theory, we study a specific yet general class of such objects: we focus on surfaces with cylindrical coordinates and their associated grasps. This yields a rich infinite dimensional space, the *Grasp*

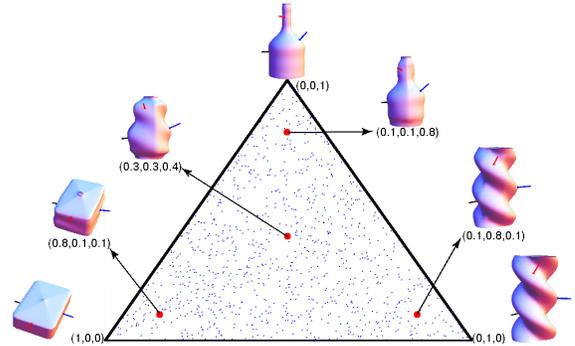


Fig. 1: Samples from a finite-dimensional subset of the Grasp Moduli Space spanned by the three surfaces displayed at the vertices of the triangle.

Moduli Space for such surfaces, which we define here and which parametrizes continuous deformations of both objects and grasps. In this work, we study some fundamental properties of this space, including a construction of finite dimensional subsets induced by any finite set of ‘shape examples’. Our definition of the Grasp Moduli Space is motivated by the belief that, given the definitions and results presented in this paper, it will be possible to define deterministic and probabilistic methods for grasping which are able to utilize knowledge of observed stable grasps on different surfaces. The contributions of this paper can be summarized as follows:

- We introduce the Grasp Moduli Space \mathcal{G}^{cyl} incorporating information about both grasps and associated surfaces. We define a metric on \mathcal{G}^{cyl} , describe continuous deformations in \mathcal{G}^{cyl} and study finite dimensional subsets of \mathcal{G}^{cyl} induced by any finite collection of such surfaces.
- We provide experiments showing the evolution of stable grasps under small deformations of the underlying surface, the contact positions of a grasp, or both, inside the Grasp Moduli Space \mathcal{G}^{cyl} .
- We investigate the change in grasp quality under large deformations of surfaces and grasp configurations in \mathcal{G}^{cyl} .
- We formulate a grasp quality gradient approach in \mathcal{G}^{cyl} for transferring stable grasps from one surface to another surface and evaluate our approach in simulation.
- We develop an efficient heuristic for finding stable grasps on surfaces of revolution corresponding to natural grasps used by humans. We then show how to continuously deform any surface with cylindrical coordinates to an associated canonical surface of revolution and use the above gradient approach to generate stable grasps based on our heuristic.

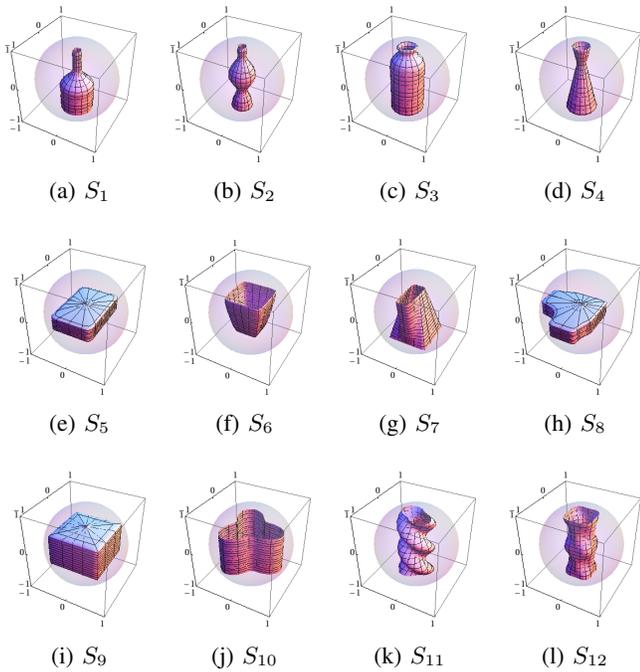


Fig. 2: Example surfaces with cylindrical coordinates. Figures a-d) display surfaces of revolution, while the surfaces in e-l) do not exhibit the rotational symmetry.

II. BACKGROUND AND RELATED WORK

The idea of continuously varying shapes has implicitly surfaced in the grasping community in works using simple primitives such as boxes and quadrics that depend continuously on their parameters. Similarly, one can consider the space of continuously varying hand configurations. A successful technique is to represent robot hand configurations as elements of a vector space and to employ principal component analysis to determine postural synergies [8, 18]. These then define a low-dimensional representation which can be used to control a robotic hand. Various extensions of this concept, such as soft synergies [3] and adaptive synergies [6] have been proposed. In another related work, [14] proposed an alternative continuous representation of grasp configurations based on simplicial complexes. These representations do however focus on continuous deformations of hand postures rather than capturing the space of both objects and contact point positions jointly. In this work, we shall concentrate on grasping at the contact level [19]. Hand kinematics do not feature in our current formulation, but can be integrated with our approach in a later planning stage which attempts to position the robot hand so that a given contact configuration is realised.

Another field where the idea of continuous deformations of objects has recently shown promise in applications is computer graphics. There, continuous representations of 3D objects can be used for example for character animation, database-search [16] and mesh-morphing [1].

Most closely related to our approach in philosophy is however the use of continuous deformations in algebraic topology [9] and algebraic geometry. In algebraic geometry in particular,

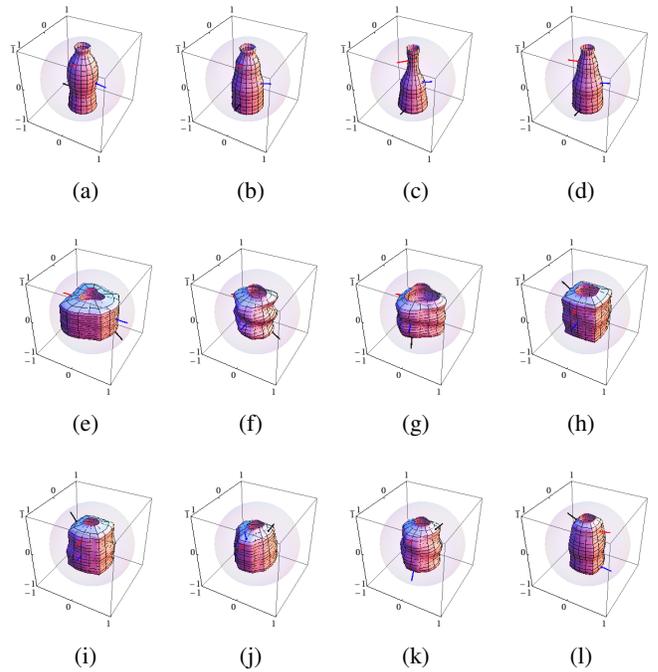


Fig. 3: Samples from the Grasp Moduli Space. Figure a-d) and e-h) show grasp/surface combinations based on the surfaces in Fig. 2a-d) and 2e-h) respectively, while the surfaces in i-l) are based on convex combinations of all the surfaces in Fig. 2.

the notion of ‘shapes up to continuous deformations’ has been used with great success for the classification of smooth manifolds.

Several grasp quality measures have been developed in order to evaluate and rank grasp hypotheses according to their physical stability [4, 7, 19, 15]. In this work, we will consider a well-established L^1 grasp quality measure Q_l introduced by [7] which we shall review now. For a given configuration $\{(c_i, n_i)\}_{i=1}^m$ of m contact positions $c_i \in \mathbb{R}^3$ and corresponding inward pointing unit normal vectors $n_i \in \mathbb{S}^2 \subset \mathbb{R}^3$ and a centre of mass $z \in \mathbb{R}^3$, let us denote this data by $g = (c_1, \dots, c_m, n_1, \dots, n_m, z) \in \mathbb{R}^{3m} \times (\mathbb{S}^2)^m \times \mathbb{R}^3$, so that g lives in a $5m + 3$ dimensional space. For contact c_i , the Coulomb friction model states that, for any force f_i applied at c_i , we have $\|f_i^t\| \leq \mu f_i^\perp$, constraining f_i to lie in the friction cone at c_i , where $f_i^\perp \in \mathbb{R}$ and $f_i^\perp n_i$ denotes the component of f_i along the normal vector n_i on the surface S containing the contact point c_i , and f_i^t denotes the component of f_i tangent to S at c_i . The friction cone $C_i = \{f_i \in \mathbb{R}^3 : \|f_i^t\| \leq \mu f_i^\perp\}$ can be approximated by $\{\sum_{j=1}^l \alpha_{ij} f_{ij} : \alpha_{ij} \geq 0\}$ for $l \geq 3$ uniformly spaced vectors $f_{ij} \in C_i$, $j \in \{1, \dots, l\}$ per cone C_i , satisfying $\langle f_{ij}, n_i \rangle = 1$. We work with $l = 8$ in this paper. The space of wrenches satisfying $\sum_{i=1}^m |f_i^\perp| \leq 1$ is then approximated by $\text{Conv}(0, S(g))$ using the convex hull $S(g) = \text{Conv}(w_{ij} : i = 1, \dots, m \text{ and } j = 1, \dots, l) \subset \mathbb{R}^6$, where $w_{ij} = (f_{ij}, (c_i - z) \times f_{ij})$. $Q_l(g)$ is defined to be the distance from the origin to the boundary of $S(g)$ if $0 \in \text{Int}(S(g))$ and is zero otherwise. In practice, $Q_l(g)$ is determined by computing a description of $S(g)$ as an intersection of affine

half-spaces, $S(g) = \bigcap_{i=1,\dots,s} \{x \in \mathbb{R}^6 : \langle x, v_i \rangle \leq \lambda_i\}$ and then $Q_l(g) = \max(0, \min_{i=1,\dots,s} \lambda_i)$. Positive values of Q_l correspond to physically stable grasps since arbitrary wrenches of magnitude less than or equal to $Q_l(g)$ can then be resisted [7]. Note that alternative grasp quality measures [7] can be defined. One could for example consider forces such that $\max_i \|f_i^\perp\| \leq 1$, but the resulting L^∞ quality measure Q'_l is computationally harder to determine and, since $Q_l(g) < Q'_l(g)$, stable grasps with respect to Q_l are also stable with respect to Q'_l , making Q_l a more desirable measure for our current investigation. For our experiments, we will consider the case of $m = 3$ contact points. While Q_l provides a physically well-motivated measure for grasp stability, it is a rather difficult function to analyse analytically due to its complicated definition involving convex hulls. In 2D, progress on the existence of stable grasps [15] and in particular on the synthesis of optimal two and three finger grasps has been made, but in 3D, state of the art approaches to grasp synthesis currently still employ various heuristics [17] combined with random sampling [5] in order to search for optimal grasp configurations.

III. OUR FRAMEWORK

Representing surface and grasp configurations

Let us observe that, for the purpose of grasping, most real world objects can be locally approximated by smooth surfaces in \mathbb{R}^3 . A large subset of these surfaces can furthermore be specified by cylindrical coordinates: $S_{f,a,b} = \{(f(u, \theta) \cos \theta, f(u, \theta) \sin \theta, (1-u)a + ub) : u \in [0, 1], \theta \in \mathbb{S}^1\}$ for some smooth bounded function $f : [0, 1] \times \mathbb{S}^1 \rightarrow \mathbb{R}_{>0}$ with bounded derivatives and where the height of the surface is bounded by $[a, b]$, with $a < b$. We call the set of these surfaces with cylindrical coordinates \mathcal{M}^{cyl} . A familiar subset of \mathcal{M}^{cyl} is the set of surfaces of revolution $\mathcal{M}^{rev} \subset \mathcal{M}^{cyl}$ given by surfaces of the form $S_{f,a,b} = \{(f(u) \cos \theta, f(u) \sin \theta, (1-u)a + ub) : u \in [0, 1], \theta \in \mathbb{S}^1\}$, where $a < b$, $f : [0, 1] \rightarrow \mathbb{R}_{>0}$ is again smooth, bounded and with bounded derivatives and independent of the angular coordinate. Fig. 2 displays some examples. Note that we can even model non-smooth shapes such as a box as elements of \mathcal{M}^{cyl} after smoothing the edges and after removing a small neighbourhood around the z -axis at the top and bottom as in Fig. 2e). Most everyday objects such as bottles, glasses, handle-parts, etc. can hence be considered either as elements of \mathcal{M}^{cyl} or can be decomposed into parts which lie in \mathcal{M}^{cyl} . While we are not investigating point-cloud data here, one can obtain coordinates for a suitable point-cloud P , e.g. by performing regression in cylindrical coordinates on \mathbb{R}^3 with respect to the most dominant axis of P . For $S_{f,a,b}, S_{g,c,d} \in \mathcal{M}^{cyl}$ and $\alpha, \beta > 0$, we define $\alpha S_{f,a,b} + \beta S_{g,c,d} = S_{\alpha f + \beta g, \alpha a + \beta c, \alpha b + \beta d}$.

Lemma III.1. \mathcal{M}^{cyl} and \mathcal{M}^{rev} are convex cones.

Proof: Let $S_{f,a,b}, S_{g,c,d} \in \mathcal{M}^{cyl}$ and $\alpha, \beta > 0$. Clearly $\alpha f + \beta g$ is smooth, positive and bounded and $\alpha a + \beta c < \alpha b + \beta d$. Hence, the result about \mathcal{M}^{cyl} follows. The derivation for \mathcal{M}^{rev} is similar. ■

Note that \mathcal{M}^{cyl} and \mathcal{M}^{rev} are infinite dimensional. Given a set $X = \{S_{f_1, a_1, b_1}, \dots, S_{f_n, a_n, b_n}\} \subset \mathcal{M}^{cyl}$ (or \mathcal{M}^{rev}), it is natural to consider their convex hull $Conv(X) = \{\sum_{i=1}^n \alpha_i S_{f_i, a_i, b_i} : \alpha_i \geq 0, \sum_{i=1}^n \alpha_i = 1\}$ which is finite-dimensional and a subset of \mathcal{M}^{cyl} (\mathcal{M}^{rev} respectively). $Conv(X)$ yields a natural bounded and finite-dimensional ‘shape-space’ spanned by the discrete ‘example set’ X of observed shapes. Since there exists natural probability distributions on any simplex (e.g. the Dirichlet distribution), it is also possible to sample random shapes from $Conv(X)$ as in Fig. 3 and to build probabilistic methods on this space. In Fig. 1, we display a projection of a 2-simplex spanned by the shapes displayed at its corners together with a few uniform random samples. Note that one could also consider the unbounded but finite-dimensional space $Cone(X) = \{\sum_{i=1}^n \alpha_i S_{f_i, a_i, b_i} : \alpha_i > 0\}$, which we shall however not study here. In this work, \mathcal{M}^{cyl} and \mathcal{M}^{rev} will be used as natural ‘shapes spaces’ for the purpose of grasping. Next, we study grasps on these surfaces and endow the resulting space with a metric.

Given a surface $S = S_{f,a,b} \in \mathcal{M}^{cyl}$ and a point $p \in S_{f,a,b}$, we can parametrize p by $p = c_S(u, \theta) = (f(u, \theta) \cos(\theta), f(u, \theta) \sin(\theta), (1-u)a + ub)$ for some $(u, \theta) \in [0, 1] \times \mathbb{S}^1$. We can furthermore compute two tangent vectors to S at p : $t_1(u, \theta) = \frac{\partial c_S}{\partial u}$, $t_2(u, \theta) = \frac{\partial c_S}{\partial \theta}$ and the corresponding unit normal vector $n_S(u, \theta) = \frac{t_1(u, \theta) \times t_2(u, \theta)}{\|t_1(u, \theta) \times t_2(u, \theta)\|}$. Note that, since f has bounded derivatives, n_S takes values in the twice-punctured sphere $\mathbb{S}_\pm^2 = \mathbb{S}^2 - \{\pm(0, 0, 1)\}$. A grasp with m contact points can now be described by $(u_1, \dots, u_m, \theta_1, \dots, \theta_m) \in [0, 1]^m \times (\mathbb{S}^1)^m = \mathcal{C}^m$ using the maps c_S and n_S . Given a uniform mass distribution on the surface, its centre of mass z_S can be computed by the formula

$$z_S = \frac{\int_0^1 \int_0^{2\pi} c_S(u, \theta) \|t_1(u, \theta) \times t_2(u, \theta)\| d\theta du}{\int_0^1 \int_0^{2\pi} \|t_1(u, \theta) \times t_2(u, \theta)\| d\theta du}.$$

The grasp quality Q_l and \mathcal{M}^{cyl}

In order to determine the grasp quality Q_l for grasps on surfaces in \mathcal{M}^{cyl} , we first need to fix our approximation of the force-cones. For the standard normal vector $n = (0, 0, 1)$, and for friction coefficient $\mu > 0$, we define the edges of the force-cone along $(0, 0, 1)$ by $f_j = (\mu \cos(\frac{2\pi j}{l}), \mu \sin(\frac{2\pi j}{l}), 1)$, for $j = 1, \dots, l$. For any other unit normal vector $n \neq \pm(0, 0, 1)$, define f_j by rotating the f_j for the normal $(0, 0, 1)$ by an angle of $\arccos(\langle (0, 0, 1), n \rangle)$ about the axis $(0, 0, 1) \times n$. If $n = (0, 0, -1)$, we rotate by the diagonal matrix with entries $(1, -1, -1)$ along the diagonal. The above assignment of friction cone edges to normal directions is clearly continuous for $n \in \mathbb{S}_\pm^2$.

Lemma III.2. Q_l is a continuous function on $\mathbb{R}^{3m} \times (\mathbb{S}_\pm^2)^m \times \mathbb{R}^3$.

Proof: The force cones depend continuously on the input data on the given domain and so does the convex hull. Finally, the distance function from the origin to the nearest facet of the convex hull is clearly continuous under continuous variations of the convex hull. ■

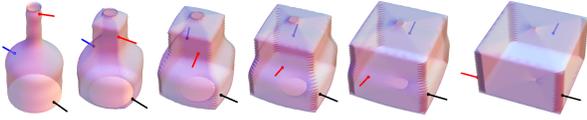


Fig. 4: Continuous combined surface and grasp deformation $\gamma : [0, 1] \rightarrow \mathcal{G}^{cyl}(3)$ between two points in $\mathcal{G}^{cyl}(3)$.

Shape and grasp deformations

In order to be able to talk about continuous grasp/surface deformations, we make the following definition:

Definition III.3. The Grasp Moduli Space for surfaces with cylindrical coordinates and grasps with m contact points $\mathcal{G}^{cyl}(m)$ is given by $\mathcal{G}^{cyl}(m) = \mathcal{M}^{cyl} \times \mathcal{C}^m$, where $\mathcal{C}^m = [0, 1]^m \times (\mathbb{S}^1)^m$. Similarly, we define the Grasp Moduli Space for surfaces of revolution with m contact points by $\mathcal{G}^{rev}(m) = \mathcal{M}^{rev} \times \mathcal{C}^m$ and note that $\mathcal{G}^{rev}(m) \subset \mathcal{G}^{cyl}(m)$.

Note that the spaces $\mathcal{G}^{rev}(m)$, $\mathcal{G}^{cyl}(m)$ are infinite dimensional since the spaces \mathcal{M}^{rev} , \mathcal{M}^{cyl} are. Since a point $(S_{f,a,b}, g) \in \mathcal{M}^{cyl} \times \mathcal{C}^m = \mathcal{G}^{cyl}(m)$ (or $\mathcal{G}^{rev}(m)$) describes both a particular surface $S_{f,a,b}$ and a grasp g on that surface, we are now able to describe various interesting continuous deformations in that space: suppose that $(S, u, \theta), (S', u', \theta') \in \mathcal{G}^{cyl}(m)$, with $u, u' \in [0, 1]^m$, and $\theta, \theta' \in (\mathbb{S}^1)^m$. Consider the curve $\gamma : [0, 1] \rightarrow \mathcal{G}^{cyl}(m)$ given by

$$\gamma(t) = ((1-t)S + tS', (1-t)u + tu', \varphi(\theta_1, \theta'_1, t), \dots, \varphi(\theta_m, \theta'_m, t)), \quad (\text{III.1})$$

where we let $t \mapsto \varphi(\theta_i, \theta'_i, t)$ be an interpolation between θ_i and θ'_i along a shortest path on the circle, i.e. we pick coordinates $\theta_i, \theta'_i \in [0, 2\pi]$ on the circle and define $\varphi(\theta_i, \theta'_i, t) = (1-t)\theta_i + t\theta'_i$ if $|\theta_i - \theta'_i| \leq \pi$ and otherwise

$$\varphi(\theta_i, \theta'_i, t) = \begin{cases} (1-t)(\theta_i + 2\pi) + t\theta'_i \pmod{2\pi}, & \theta_i < \theta'_i \\ (1-t)\theta_i + t(\theta'_i + 2\pi) \pmod{2\pi}, & \theta_i > \theta'_i. \end{cases}$$

We can now distinguish three interesting types of deformations in $\mathcal{G}^{cyl}(m)$. *Surface-only deformations*, where $u = u', \theta = \theta'$, so that the cylindrical coordinates of the contact points stay constant, but $S \neq S'$, *grasp-only deformations*, where $S = S'$, but we vary the contact points, and *combined deformations*, where both the surface and the grasp are being deformed. Note that the simple curve γ described above is only one example (as in Fig. 4) of how to continuously interpolate between two grasp/surface configurations. Optimal curves, e.g. under task-constraints could be another interesting future line of research. We now define a simple metric on the space of grasps. Let $g, g' \in \mathbb{R}^{3m} \times (\mathbb{S}^2)^m \times \mathbb{R}^3$, with

$$g = (c_i, \dots, c_m, n_1, \dots, n_m, z), \\ g' = (c'_i, \dots, c'_m, n'_1, \dots, n'_m, z').$$

We define

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

where $d_{\mathbb{R}^3}(x, y) = \|x - y\|$ is the standard Euclidean metric and we think of n_i, n'_i as embedded in \mathbb{R}^3 , in order to measure distances between normal vectors.

So far, we have defined a metric on the space of grasps without any notion of a surface. We now consider a distance on our shape space \mathcal{M}^{cyl} which can easily be related to our distance on grasp configurations. For $S, S' \in \mathcal{M}^{cyl}$,

$$d(S, S') = \sup_{h \in [0, 1], \theta \in \mathbb{S}^1} \max \{ d_{\mathbb{R}^3}(c_S(u, \theta), c_{S'}(u, \theta)), d_{\mathbb{R}^3}(n_S(u, \theta), n_{S'}(u, \theta)) \}.$$

Note that, thinking of \mathbb{S}^1 as the unit circle in \mathbb{R}^2 , we can endow \mathbb{S}^1 with the simple metric $d_{\mathbb{S}^1}(x, y) = \|x - y\|$, where $\|\cdot\|$ denotes the Euclidean distance in \mathbb{R}^2 (alternatively, we could measure differences in angles). Similarly, on $I = [0, 1]$ $d_I(x, y) = |x - y|$ provides us with a distance measure. Summing these distances and the above metric d component-wise, we can hence obtain a metric on $\mathcal{G}^{cyl}(m) = \mathcal{M}^{cyl} \times \mathcal{C}^m$, which can be used to quantify distances between grasp/surface combinations.

Definition III.4. Let $S_1 = S_{f_1, a_1, b_1}, S_2 = S_{f_2, a_2, b_2}$ be surfaces with cylindrical coordinates and let

$$g = (c_{S_1}(u_1, \theta_1), \dots, c_{S_1}(u_m, \theta_m), n_{S_1}(u_1, \theta_1), \dots, n_{S_1}(u_m, \theta_m), z_{S_1})$$

be a grasp configuration on S_{f_1, a_1, b_1} . The *naïve transfer* of g to S_{f_2, a_2, b_2} is the grasp on S_{f_2, a_2, b_2} given by

$$g' = (c_{S_2}(u_1, \theta_1), \dots, c_{S_2}(u_m, \theta_m), n_{S_2}(u_1, \theta_1), \dots, n_{S_2}(u_m, \theta_m), z_{S_2}).$$

Lemma III.5 (Transfer lemma). *Suppose two surfaces with cylindrical coordinates S_1, S_2 satisfy $d(S_1, S_2) \leq \varepsilon$ for some $\varepsilon \geq 0$ and let*

$$g = (c_{S_1}(u_1, \theta_1), \dots, c_{S_1}(u_m, \theta_m), n_{S_1}(u_1, \theta_1), \dots, n_{S_1}(u_m, \theta_m), z_{S_1})$$

be a grasp configuration on S_1 . Then the naïve transfer g' of g to S_2 satisfies $d(g, g') \leq \varepsilon$.

Proof: This follows directly from the definition of the distances. ■

Having related distances of shapes to distances of grasps, we now come to the connection with grasp quality. Let us consider grasps g such that the contacts c_i and the centre of mass z are constrained to lie within the ball $\mathbb{B}(r) = \{x \in \mathbb{R}^3 : \|x\| \leq r\}$. The resulting set $\mathcal{D}(r) = \mathbb{B}(r)^m \times (\mathbb{S}^2)^m \times \mathbb{B}(r)$ is bounded and we can consider, for $0 \leq \delta_- < \delta_+$, the quantity

$$\varepsilon_l(\delta_-, \delta_+, r) = \sup \{ \varepsilon : Q_l(g) > \delta_- \text{ for all } g \text{ such that } d(g, g') \leq \varepsilon \text{ for some } g' \in \mathcal{D}(r) \text{ with } Q(g') \geq \delta_+ \}.$$

This quantity, which could be studied analytically or - as we shall do - numerically, can be used to give guarantees on the existence of stable grasps near a known stable grasp

configuration. This applies both to the case where we would like to find a nearby grasp on a single given surface, or when we want to transfer a grasp to a nearby position on a *nearby surface*. We can hence already see the usefulness of specifying a single space such as the Grasp Moduli Space, since we are now able to precisely formulate questions such as ‘when can a stable grasp on surface X be transformed to a stable grasp on surface Y ’ as problems in $\mathcal{G}^{cyl}(m)$. We are especially excited about the potential future prospect of using non-parametric Bayesian methods such as Gaussian Processes on $\mathcal{G}^{cyl}(m)$, which could provide a step towards a continuous non-parametric Bayesian approach to grasp synthesis.

IV. EXPERIMENTS

For our experiments, we approximate friction cones by a convex hull of $l = 8$ uniformly spaced edges. We are interested only in grasps with three contact points lying in the bounded domain $\mathcal{D}(1)$. All grasps on convex combinations of the surfaces depicted in Fig. 2 lie in this set. We work with a friction coefficient of $\mu = 1$ throughout.

A. Local safe grasp deformation ε -neighbourhoods

In this section, we experimentally determine a lower bound ε for the quantities $\varepsilon_8(0, 0.05, 1)$ and $\varepsilon_8(0, 0.20, 1)$. This provides us with a first numerical measure for how much the quality Q_8 of a grasp varies with reasonably small changes in grasp configurations. This stage of our analysis is completely generic and does not require any notion of an object shape yet.

To determine a lower bound, we sample 100 million points \mathcal{S} from the uniform distribution on the bounded set $\mathcal{D}(1)$, for $m = 3$ contact points. This results in 83.9% unstable grasps $\mathcal{D}_0 = \{g \in \mathcal{S} : Q_8(g) = 0\}$. 4.1% of all grasps lied in the set $\mathcal{D}_{0.05} = \{g \in \mathcal{S} : Q_8(g) \geq 0.05\}$, while approximately 0.064% lied in the set $\mathcal{D}_{0.2} = \{g \in \mathcal{S} : Q_8(g) \geq 0.2\}$. The stable grasps were distributed as displayed in Fig. 5 and are increasingly concentrated towards the origin. The maximal grasp quality encountered was approximately 0.4013. Since

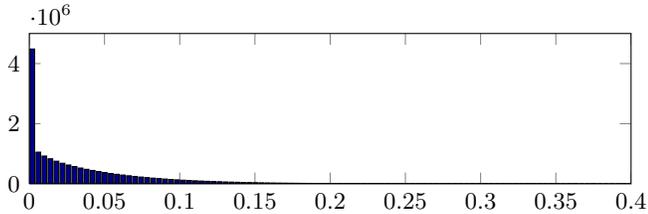


Fig. 5: Grasp quality histogram for $Q_8 > 0$, as determined from 100 million random grasps in $\mathcal{D}(1)$.

our grasp set \mathcal{S} is rather large, computing a straight-forward estimator e.g. for $\varepsilon_8(0, 0.05, 1)$ by $\varepsilon = \min_{g \in \mathcal{D}_0, g' \in \mathcal{D}_{0.05}} d(g, g')$ is computationally too expensive. Since, for any $g \in \mathcal{D}(1)$, we can however efficiently sample new points from the uniform distribution on the neighbourhood $\mathcal{N}_\varepsilon(g) = \{g' : d(g', g) \leq \varepsilon\}$, we can sample new points from the uniform distribution

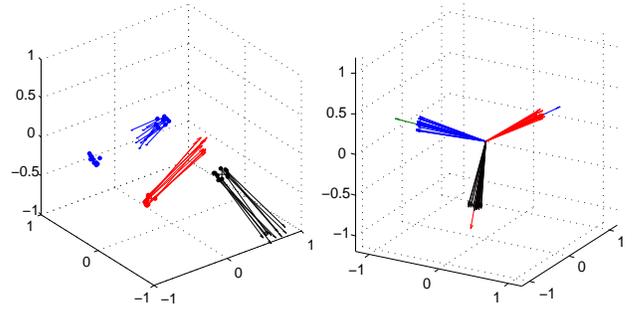


Fig. 6: On the left, we display resampled grasp configurations in a neighbourhood of size $\varepsilon = 0.1$ around a particular grasp g . Inward pointing resampled normal vectors based at the resampled contact point positions are displayed together with the resampled centre of mass. On the right, normals of a grasp g resampled in a neighbourhood of size $\varepsilon = 0.1$ are displayed separately and are based at the origin. For comparison, the original normals are drawn slightly longer than the samples.

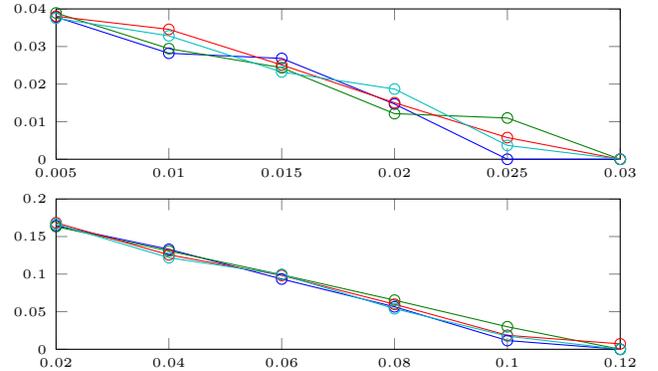


Fig. 7: The top figure displays the smallest grasp quality (vertical axis) encountered when resampling 25 grasps in a neighbourhood of size ε (horizontal axis) around approximately 41000 grasps from $\mathcal{D}_{0.05}$. The bottom figure displays the smallest grasp quality encountered when resampling 150 points in a neighbourhood of size ε (horizontal axis) for each of the approximately 64000 grasps in $\mathcal{D}_{0.2}$. The experiments are repeated four times resulting in the displayed small variations.

on $\mathcal{N}_\varepsilon(g)$ for each $g \in \mathcal{D}_{0.05}$ and test if we encounter unstable configurations for various ε settings.

For $\varepsilon_8(0, 0.05, 1)$, we worked with approximately 41000 samples from $\mathcal{D}_{0.05}$ and resampled 25 points from $\mathcal{N}_\varepsilon(g)$ for each such sample g and for $\varepsilon \in \{0.005, 0.01, 0.015, 0.02, 0.025, 0.03, 0.03\}$. Similarly, for $\varepsilon_8(0, 0.2, 1)$, we worked with all available samples in $\mathcal{D}_{0.2}$ (about 64000) and resampled 150 new points for each such sample (see Fig. 6) and for $\varepsilon \in \{0.02, 0.04, 0.06, 0.08, 0.1, 0.12\}$. To test for stability, we repeated these experiments 4 times. Fig. 7 summarizes the result. We found that the minimal grasp quality for each parameter setting remained rather stable between iterations of our resampling procedure. Our results show in particular that 0.02 and 0.1 can be used as a lower-bound estimate for $\varepsilon_8(0, 0.05, 1)$ and $\varepsilon_8(0, 0.2, 1)$ respectively. Let us remark here that the determination of optimal approximations to

$\varepsilon_l(\delta_-, \delta_+, r)$ is a very interesting topic, which we believe should be investigated in more detail in future. Importantly, estimates of this type do only need to be computed *once* for various parameter settings for later use - *e.g.* using a cluster computer. Once such constants are determined, sampling algorithms could then use this information to explore the space of possible contact point configurations more rapidly than uninformed uniform sampling.

B. An example of ε -stability under deformations in \mathcal{G}^{cyl}

We now apply the local stability result just obtained to three types of deformations in the space $\mathcal{G}^{cyl}(3)$. We consider the surface S with coordinates $(0.3, 0.3, 0.4)$ displayed in Fig. 1. We determined 200 uniformly distributed random contact positions on S using the stratified uniform sampling method on smooth embedded surfaces as explained in [2]. This involves an inversion of a strictly monotonic function on $[0, 1]$ for which we use the bisection method with 30 steps and integrations in $1D$ and $2D$, for which we use Simpson's method partitioning the domain of integration into a grid with 60 uniform steps along each dimension. We then rank all combinations of three such contact points by grasp quality. This results in $\binom{200}{3} = 1313400$ evaluations of Q_8 and takes about 31 minutes on a current computer, while the random sampling takes about half a minute. The best grasp g (displayed in blue in Fig.8) satisfied $Q_8(g) \approx 0.33$. We now experimented with using the bound on $\varepsilon_8(0, 0.2, 1)$ by deforming the grasp and surface. One might imagine that the grasp g has been previously learned by the robot and is stored in a database. Suppose now that the robot encounters the following three cases: a) the object is encountered again, but due to imprecision in control is grasped slightly differently (grasp-only deformation), b) another object S' very similar to S is encountered and the naïve transfer grasp is executed precisely (surface-only deformation), c) an object S' similar to S is encountered but the transfer grasp is executed only up to some imprecision (combined deformation). Due to our bound of the previous section, we know that any grasp g' such that $d(g, g') \leq 0.1$ should be stable and that the naïve transfer of the grasp g to any surface S' such that $d(S, S') \leq 0.1$ will also be stable by the naïve transfer lemma. We explicitly verified each of the three cases by sampling 3 random surfaces near the surface S in the simplex depicted in Fig. 1 and by sampling nearby grasps for the case a) and c). Fig. 8 depicts an explicit example, where we now allow larger deformations such that $d(g, g') \leq 0.3$ and $d(S, S') \leq 0.3$ respectively. Even for those deformations, the grasp quality of g' remained larger than 0.19 in this example.

C. Behaviour under global deformations in \mathcal{G}^{cyl}

Next, let us get a qualitative understanding of how grasp quality changes for larger deformations along a continuous path in the Grasp Moduli Space $\mathcal{G}^{cyl}(3)$. To obtain a starting set of interesting grasps, we uniformly sample 200 contact positions on each of our model surfaces displayed in Fig. 2 using the uniform sampling procedure just described. We again

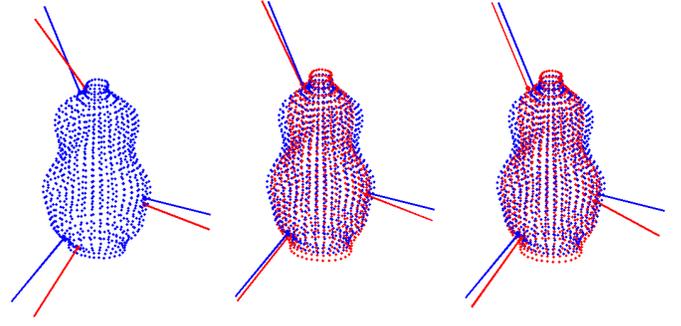


Fig. 8: Surface S with a stable grasp g (blue lines) and three deformations in the Grasp Moduli Space in a ball of radius 0.3, where only the contact and only the surface are deformed by a small amount in a) and b) respectively and where both are deformed in c). In all cases, the grasp remains stable after deformation.

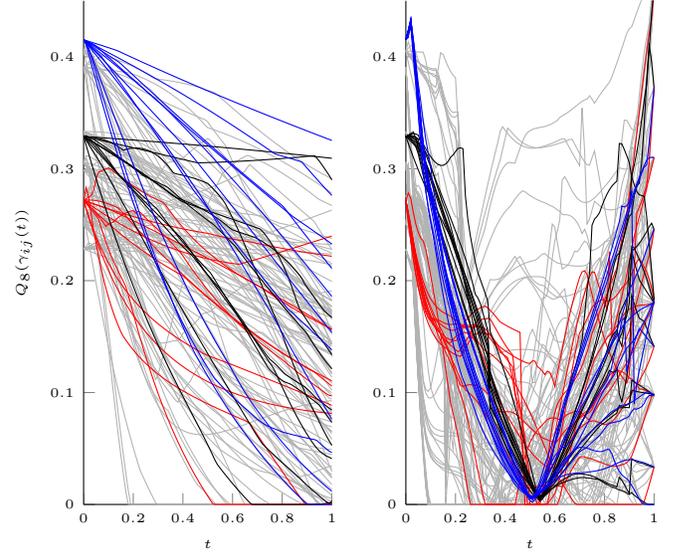


Fig. 9: Plots of $Q_8(\gamma_{ij}(t))$ against $t \in [0, 1]$, where $\gamma_{ij}(t)$ denotes the naïve transfer of a fixed grasp $g_i \in S_i$ to S_j , $i \neq j$ on the left and the interpolation transfer from $g_i \in S_i$ to the grasp with coordinates $u = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, $\theta = (0, \frac{2\pi}{3}, \frac{4\pi}{3})$ on S_j , $i \neq j$, on the right. The red, black and blue lines correspond to starting from the surfaces displayed in Fig. 2a, i) and j) respectively.

calculate the grasp quality of all triples of such contacts up to permutation and use the best grasp g_i on S_i , for $i = 1, \dots, 12$ as a model grasp on each surface. The left part of Fig. 9 displays the variation in grasp quality when we try to apply the naïve transfer of the grasp g_i from S_i to S_j , $i \neq j$, by simply keeping the same cylindrical coordinates of the grasp while deforming the underlying surfaces via $\gamma_{ij}(t) = (1 - t)S_i + tS_j$ for $t \in [0, 1]$. The right part of that figure shows the interpolation transfer between $g_i \in S_i$ and the grasp on surface S_j with coordinates $u = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, $\theta = (0, \frac{2\pi}{3}, \frac{4\pi}{3})$ and using $\gamma(t)$ as in Eq. III.1. As can be seen from the figure, Q_8 behaves in fact rather tamely in our naïve transfer examples. While Q_8 is not a smooth function, it seems to be Lipschitz continuous - as already indicated by our experiment a). For the interpolation transfer, we note a higher amount of variability in our example

and that a large amount of target grasps with coordinates given by (u, θ) are in fact stable at the end of the deformation. Let us also remark here that there are many interesting open questions about interpolation transfers that can be asked, such how to characterize subsets of grasp/surface configurations in $\mathcal{G}^{cyl}(3)$ for which such transfers lead to stable grasps.

D. Gradient methods on \mathcal{G}^{cyl}

Given the ability to apply infinitesimal deformations to grasps and surfaces, variational and gradient based methods immediately come to mind. In the case of a fixed surface, a gradient ascent with respect to grasp quality can for example be considered. To the best of our knowledge, this idea has only been investigated in the simple case of a sphere [11]. Since we are however also able to deform the surface, we focus on the interesting question of how to optimally transfer a stable grasp $g \in S_i$ to the surface S_j while maintaining stability and where S_1, \dots, S_{12} again denote the surfaces depicted in Fig. 2. We propose to use a sub-division $\{t_0, \dots, t_{N_1}\} \subset [0, 1]$ with $t_s = \frac{s}{N_1}$ and $N_1 = 100$ and consider the curves of surfaces $\alpha_{ij}(t) = (1-t)S_i + tS_j$. While Q_8 is not differentiable everywhere, we none-the-less evaluate a gradient descent inspired approach on $-Q_8$ by approximating $\frac{\partial Q_8}{\partial x_i} \approx \frac{Q_8(x+he_i) - Q_8(x)}{h}$, where e_i denotes the i -th standard basis vector. We chose $h = 10^{-5}$ in our experiments. We now propose the following simple algorithm in Grasp Moduli Space, transferring a grasp $g \in S$ to a stable grasp on S' : In step s , we iterate $N_2 = 20$ gradient descent steps, with $x_{n+1} = x_n + \delta \nabla Q(x_n)$ and step length $\delta = 0.001$ on the surface $\alpha_{ij}(t_s)$ to reach a local maximum on $\alpha_{ij}(t_s)$. We then apply the naïve transfer to the nearby surface $\alpha_{ij}(t_{s+1})$ and proceed to step $s + 1$ until $\alpha_{ij}(1) = S_j$ is reached. In order to measure the improvement of this approach relative to the global naïve transfer considered in the previous experiment, we apply our gradient method to the grasps $g_i \in S_i$ considered previously. We found that our method typically improves the quality of the grasp during the deformation when compared to the naïve method, but the absolute grasp quality can undergo non-monotonic variations during the deformation. Fig. 10 shows the final relative improvement in grasp quality. As we can see, our approach typically improves the grasp quality substantially. Furthermore, all the resulting grasps on the target surfaces were stable in our experiment. Let us remark here also that we did not exclude the case of a transfer ‘from S_i to S_i ’, which corresponds to performing gradient based grasp optimization on a fixed surface.

E. The role of $\mathcal{M}^{rev} \subset \mathcal{M}^{cyl}$

Recall that we can think of surfaces of revolution as forming a natural infinite-dimensional subset $\mathcal{M}^{rev} \subset \mathcal{M}^{cyl}$. When we determined best grasps on the surfaces of revolution displayed in Fig. 2a-d) by brute-force search for experiment c), we observed that the best grasps on surfaces of revolution seemed to occur along a single plane containing the axis of rotation, which corresponds to a natural grasp used by humans as well. We hence propose exploiting the rotational symmetry of these

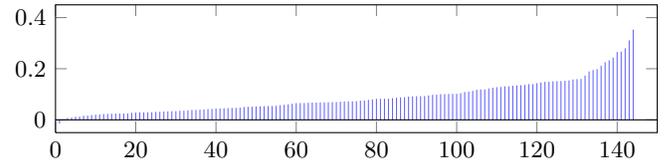


Fig. 10: Sorted differences in final grasp quality between our gradient approach and the naïve transfer approach when transferring from S_i to S_j for all $12^2 = 144$ combinations of surfaces and starting from our stable grasps on S_i (note, we include the case $i = j$ here). In all but one case, the final grasp quality using the gradient descent is higher. Remarkably, all final grasps, even those grasps that were unstable under the naïve transfer (see $t = 1$ on the left part of Fig. 9) are stable on the target surface when the gradient method is used.

surfaces and suggest the following heuristic for determining stable grasps on such surfaces:

- Let $U = \{0, \frac{1}{m}, \dots, \frac{m-1}{m}, 1\}$, for some $m \in \mathbb{N}$.
- Compute grasp quality Q_8 for the triples of contact points with cylindrical coordinates $(u_1, 0), (u_2, \pi), (u_3, \pi) \in [0, 1] \times \mathbb{S}^1$, where $u_1, u_2, u_3 \in U$ and $u_2 < u_3$.
- Select a grasp with optimal grasp quality Q_8 from the above triples as a grasp candidate.

To test our heuristic, we considered the finite dimensional subset of $\mathcal{G}^{cyl}(3)$ spanned by the four surfaces of revolution displayed in Fig. 2a-d) and sampled 100 uniformly distributed convex combinations of those surfaces and 100 uniformly distributed contact points on each such surface using stratified random sampling. We then computed the grasp quality for every three-element subset of the contact points on each of the 100 surfaces and stored the best grasp for each surface resulting from this brute-force method. This computation took about 250 seconds per surface, resulted in $\binom{100}{3} = 161700$ evaluations of Q_8 per surface and yielded stable grasps distributed as displayed in Fig. 11 and with a mean grasp quality of 0.2435. We then applied our heuristic with only $m = 10$ height sub-divisions, which took about 0.89 seconds per surface and which required just $10 \binom{10}{2} = 450$ evaluations of Q_8 per surface. Again, all resulting grasps had positive grasp quality distributed as in Fig. 11. Interestingly, the mean quality obtained using our heuristic was 0.2521 and hence *higher* than the mean quality obtained using the time-consuming brute force method.

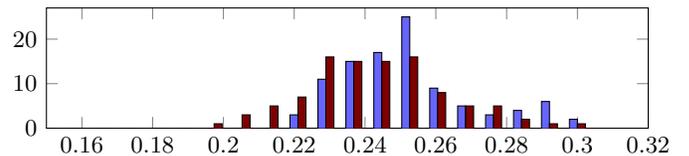


Fig. 11: Quality histogram for grasp hypotheses over 100 random surfaces of revolution spanned by the examples in Fig. 2a-d) using brute-force evaluation on 100 contact points per surface (red) and using our heuristic (blue).

Note that our simple heuristic now gives us an example for how we can reason efficiently about a subset (e.g. $\mathcal{G}^{rev}(3)$) of the Grasp Moduli Space $\mathcal{G}^{cyl}(3)$. It seems hence natural to be

able to apply this strategy for surfaces that are close to surfaces of revolution in our Grasp Moduli Space. Furthermore, there exists a canonical projection $\pi : \mathcal{M}^{cyl} \rightarrow \mathcal{M}^{rev}$ defined as follows: Given a surface $S_{f,a,b}$, we define $\pi(S_{f,a,b}) = S_{\hat{f},a,b}$, where $\hat{f}(u) = \frac{1}{2\pi} \int_0^{2\pi} f(u, \theta) d\theta$. We can continuously move towards this projection using the curve $\gamma(t) = (1-t)S + t\pi(S)$.

Let us now consider the following deformation based grasp synthesis strategy: for any surface $S \in \mathcal{M}^{cyl}$, we determine an optimal grasp g on $\pi(S)$ using our heuristic and then apply our gradient method to transfer g back to a grasp on S . Fig. 12 displays the differences in grasp quality when compared to the brute-force search for best grasps on the given surfaces using 200 uniform random samples per surface. Our method can identify stable grasps in these cases, while requiring only about 4 seconds for each surface when we chose $N_1 = N_2 = 10$, $\delta = 0.01$, $h = 10^{-5}$ and $m = 10$ for our gradient descent and heuristic parameters and when approximating centres of mass during the deformation using Simpson's integral formula on a 20 by 20 grid. As we can see, the grasp qualities are all positive but on average not as high as using brute force search. However, our method took only 4 seconds per surface as compared to 31 minutes using brute force search.

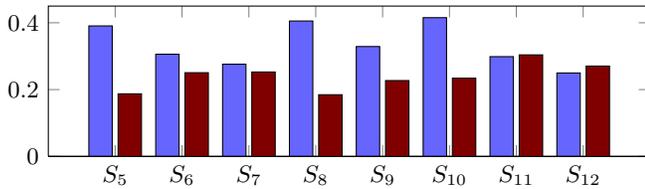


Fig. 12: Comparison of grasp quality obtained using brute-force search on 200 random contact points on the surfaces in Fig. 2e-l) (blue) and using our deformation to a surface of revolution approach (red).

V. CONCLUSION AND FUTURE WORK

In this work, we introduce a new view on contact based grasp analysis based on a single infinite dimensional space – the Grasp Moduli Space – in which continuous deformations of surface and grasp configurations can be described by continuous curves. We study finite-dimensional subsets of this space induced by data as well as the local stability of a grasp quality function. A new grasp synthesis approach based object/grasp deformations to a canonical surface of revolution is discussed as well as a gradient based approach for transferring and interpolating between grasps on differing surfaces. We believe that our framework could in future be used to develop novel probabilistic as well as optimization based techniques for grasping based on the concept of continuous deformations of both objects and grasps. While our method is currently based on contact-level information, we believe that the inclusion of robot hand kinematics in our framework could provide another promising future research direction.

VI. ACKNOWLEDGEMENTS

This work was supported by the Swedish Foundation for Strategic Research and the EU grants FLEXBOT (FP7-ERC-279933) and TOMSY (IST-FP7-270436).

REFERENCES

- [1] M. Alexa. Recent advances in mesh morphing. *Computer Graphics Forum*, 21:173–196, 2002.
- [2] J. Arvo. Stratified sampling of 2-manifolds. *SIGGRAPH 2001 Course Notes*, 29:2, 2001.
- [3] A. Bicchi, M. Gabbicini, and M. Santello. Modelling natural and artificial hands with synergies. *Phil. Trans. R. Soc. B*, 366(1581):3153–3161, 2011.
- [4] A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *IEEE ICRA*, pages 348–353 vol.1, 2000.
- [5] C. Borst, M. Fischer, and G. Hirzinger. Grasping the dice by dicing the grasp. In *IEEE/RSJ IROS*, pages 3692–3697 vol.4, 2003.
- [6] M.G. Catalano, G. Grioli, A. Serio, E. Farnioli, C. Piazza, and A. Bicchi. Adaptive synergies for a humanoid robot hand. In *IEEE–RAS Int. Conf. on Hum. Robots*, 2012.
- [7] C. Ferrari and J. Canny. Planning optimal grasps. In *IEEE ICRA*, pages 2290–2295 vol.3, 1992.
- [8] M. Gabbicini, A. Bicchi, D. Prattichizzo, and M. Malvezzi. On the role of hand synergies in the optimal choice of grasping forces. *Autonomous Robots*, 31(2-3):235–252, 2011.
- [9] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [10] K. Huebner, S. Ruthotto, and D. Kragic. Minimum volume bounding box decomposition for shape approximation in robot grasping. In *IEEE ICRA*, pages 1628–1633, 2008.
- [11] G. Liu, J. Xu, X. Wang, and Z. Li. On Quality Functions for Grasp Synthesis, Fixture Planning, and Coordinated Manipulation. *IEEE Trans. on Aut. Sci. and Eng.*, 2004.
- [12] M. Madry, D. Song, and D. Kragic. From object categories to grasp transfer using probabilistic reasoning. In *IEEE ICRA*, pages 1716–1723, 2012.
- [13] A.T. Miller, S. Knoop, H.I. Christensen, and P.K. Allen. Automatic grasp planning using shape primitives. In *IEEE ICRA*, pages 1824–1829 vol.2, 2003.
- [14] T.N. Nguyen and H.E. Stephanou. A topological algorithm for continuous grasp planning. In *IEEE ICRA*, pages 670–675 vol.1, 1990.
- [15] V.-D. Nguyen. *The Synthesis of Stable Force-Closure Grasps*. 1986.
- [16] M. Ovsjanikov, W. Li, L. Guibas, and N.J. Mitra. Exploration of continuous variability in collections of 3D shapes. In *ACM SIGGRAPH*, pages 33:1–33:10, 2011.
- [17] 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*, pages 1781–1788, 2011.
- [18] J. Romero, T. Feix, H. Kjellström, and D. Kragic. Spatio-temporal modeling of grasping actions. In *IEEE/RSJ IROS*, pages 2103–2108, 2010.
- [19] A. Sahbani, S. El-Khoury, and P. Bidaud. An overview of 3D object grasp synthesis algorithms. *Robotics and Autonomous Systems*, 60(3):326–336, 2012.