GP-GPIS-OPT: Grasp Planning Under Shape Uncertainty Using Gaussian Process Implicit Surfaces and Sequential Convex Programming

Jeffrey Mahler[1], Sachin Patil[1], Ben Kehoe[2], Jur van den Berg[3], Matei Ciocarlie[4], Pieter Abbeel[1], Ken Goldberg[1]

***Abstract*— Computing grasps for an object is challenging when the object geometry is not known precisely; especially for objects that are difficult for robots to perceive, such as those with specular or transparent surfaces. These properties introduce uncertainty in object geometry, but commonly used polygonal mesh-based models cannot easily be extended to represent this uncertainty. In this paper, we explore the use of Gaussian process implicit surfaces (GPISs) to represent shape uncertainty directly from RGBD point cloud observations of objects. We study the use of GPIS representations to select grasps on previously unknown objects, measuring grasp quality by the probability of force closure. Our main contribution is GP-GPIS-OPT, an algorithm for computing grasps for parallel-jaw grippers on 2D GPIS object representations. Specifically, we optimize an approximation to this quality subject to antipodal constraints on the parallel jaws using Sequential Convex Programming (SCP). We also introduce a method for visualizing 2D GPIS models based on blending shape samples from a GPIS. We test the algorithm on a set of perspective projections of objects that are difficult for robots to perceive. Our experiments suggest that GP-GPIS-OPT selects grasps with higher quality than a planner that ignores shape uncertainty on 7 of 8 of our test objects and is approximately $7.9\times$ faster than the most common existing method for grasp planning under shape uncertainty. We also test our method with physical experiments on the Willow Garage PR2 robot.**

I. Introduction

Consider the task of rapidly decluttering, packing, or unpacking objects with a robot in a warehouse or home environment, such as the objects illustrated in Fig. 1. The robot must be able to quickly compute grasps on each object in order to succeed, but this may be difficult when the geometry of objects in the environment is not known a priori. This can be further complicated by sensor noise, partial visibility, and surface properties of objects such as specularity and transparency, which lead to uncertainty about the geometry of objects. Therefore, a representation of this shape uncertainty and a method for computing stable grasps with respect to this uncertainty may be desirable for reliable grasp execution.

[1]Department of EECS; {jmahler, sachinpatil, pabbeel, goldberg}@berkeley.edu
[2]Department of ME; benk@berkeley.edu
[1−2] University of California, Berkeley, CA, USA
[3]Google Inc., CA, USA jurvandenberg@gmail.com
[4]Department of CS, Colombia University, NY, USA cmatei@cs.columbia.edu

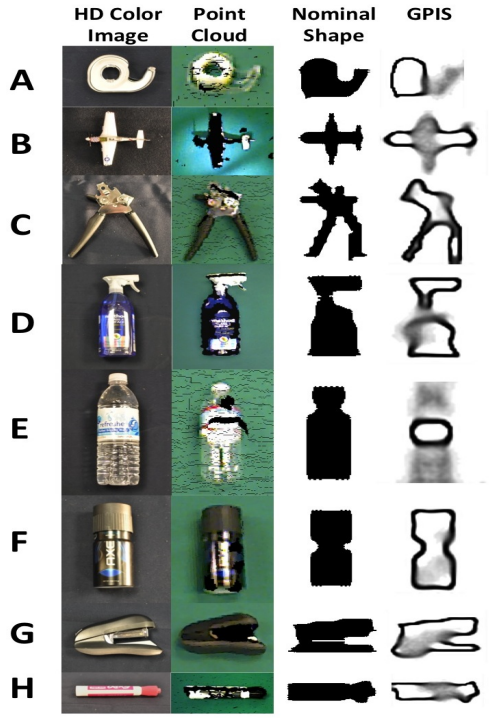

Fig. 1. Eight example objects illustrating transparency, specularity, and sensor noise: (from top to bottom) (A) tape, (B) plane, (C) can opener, (D) squirt bottle, (D) water bottle, (E) deodorant, (F) stapler, and (G) marker. Displayed from left to right are the HD color image, a point cloud observation from a Primesense Carmine, the nominal shape on a 25×25 grid, and a visualization of the GPIS representation. The GPIS visualization is blurry in areas of larger uncertainty and sharper in areas of low uncertainty. Black areas in the point cloud indicate missing observations, and in transparent areas the sensor sees through the objects and measures the depth of the table.

One promising object representation is the signed distance function (SDF), which is zero-valued at the object surface, positive-valued outside the object, and negative-valued in the object interior. This representation has been used successfully in a variety of robotic applications, such as 3D modeling, [10], [30], modeling deformations between 3D objects [39] and as potential fields for controlling grasps [11]. In this paper, we consider representing objects with a Bayesian representation of SDFs known as the Gaussian Process implicit surface (GPIS) [3], [11], [17], [33], [41] and study the use of GPIS representations to plan grasps in the presence of shape uncertainty. The key idea behind

GPIS-based representations is to represent shape uncertainty as a distribution over all possible SDFs that could fit sensor measurements of a shape. GPIS-based models have larger variances in regions where sensor measurements are missing or unreliable and lower variances where sensor measurements are known to be accurate [41].

Our main contribution in this paper is GP-GPIS-OPT, an algorithm to plan grasps for parallel-jaw grippers with high grasp quality, where quality is measured as probability of force closure given both (a) shape uncertainty represented as a GPIS and (b) known noise in physical grasp execution. We evaluate the algorithm in simulation and on the Willow Garage PR2 [1] on a set of objects are challenging for robots to perceive with visual sensors. Our experiments suggest that GP-GPIS-OPT finds grasps with higher quality than those planned without considering shape uncertainty and that GP-GPIS-OPT is approximately $7.9\times$ faster than ranking a set of random grasps by quality using Monte-Carlo integration, the most common existing method for grasp planning under shape uncertainty [7], [20], [21].

II. Related Work

Polygonal mesh models are the most common representation in robotics, but extensions of mesh models to include uncertainty is difficult. Solutions include assuming independent noise on the vertex locations, which ignores spatial noise correlations [21]. Gaussian process implicit surfaces (GPISs) are an alternative that represent objects as distribution over signed distance functions (SDFs) [10], [33], [41], and have shown promise in a number of robotic applications. Dragiev et al. [11] considered the use of GPIS for shape representations in grasping, using the maximum likelihood SDF as a potential field for a feedback controller to reach a predefined target grasp. However, the authors did not utilize uncertainty from the GPIS in their controller or to plan the target grasps. Hollinger et al. [17] created high resolution GPIS models of ship hulls to guide exploration by an underwater robot to locations with the highest shape uncertainty. Other works have used GPIS to fuse uncertain data from multiple sensors and to guide haptic exploration of object shapes [3], [19]. In comparison, our work considers how to use a GPIS representation of shape uncertainty to plan grasps with high probability of force closure.

Choosing grasps given a surface representation of an object typically focuses on finding grasps given an exact object shape by maximizing a grasp quality metric [6], [9]. A common quality measure is the Ferrari-Canny metric, which measures the ability to resist force perturbations [12] and is widely used in grasp software packages such as GraspIt! [27] and OpenGrasp [25]. Past work on grasping with uncertainty has focused on state uncertainty [13], [14], uncertainty in object pose [8], [40], [22] or uncertainty in the location of contact with the object[34].

Several works have studied the effects of shape uncertainty on grasp quality. Christopoulos et al. [7] sampled spline fits for 2-dimensional planar objects to measure the quality of potential grasps under shape uncertainty and used this to rank a set of randomly generated grasps. Kehoe et al. [20], [21] showed that adaptive sampling could be used to select grasps robust to part tolerance for parallel-jaw grippers on extruded polygon shape models. Hsaio et al. [18] studied a Bayesian framework to evaluate the probability of grasp success given uncertainty in shape and pose by simulating grasps on deterministic mesh and point cloud models. Laaksonen et al. [2], [23] used Gaussian Processes (GPs) to model distributions on grasp stability online from tactile measurements, and selected grasps using MCMC sampling. Panahi et al. [31] presented an algorithm for orienting 2D industrial parts with shape uncertainty by computing upper and lower bounds on the orientation angle of a part.

Past work on grasp optimization primarily focuses on optimizing contact points on a surface with respect to a grasp quality measure when the surface is known exactly [26], [35]. Chen et al. [5] optimized antipodal grasps to satisfy force closure for parametric surfaces without uncertainty. Ciocarlie et al. use simulated annealing to find grasps with a high Ferrari-Canny metric for the GraspIt! software, using a penalty on signed distance to the object surface to force contact with the object [9], [27]. Simulated annealing is based on random exploration of the input state space, which can avoid local minima but can be less computationally efficient than gradient-based methods [9]. Our approach can be seen as a local optimization similar to the method of Ciocarlie et al. [9] or Pokorny et al. [32] with constraints similar to those used by Chen et al. [5], but our method uses an explicit model of shape uncertainty.

III. Gaussian Process Implicit Surfaces

In this section we review Gaussian Process implicit surfaces (GPISs) and describe our method to visualize them. A signed distance field (SDF) [10] describes the shape of an object by storing the signed distance from every point in space to the nearest point on the surface. SDFs are defined as a real-valued function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $f(\mathbf{x}) > 0$ outside the object, $f(\mathbf{x}) = 0$ on the object surface, and $f(\mathbf{x}) < 0$ inside the object. The notion of SDFs can be extended to incorporate uncertainty by using Gaussian Process regression (GPR) to estimate a mean and variance over possible SDFs that fit noisy observations. This augmented definition is called a Gaussian Process implicit surface [41].

A. Gaussian Process Regression (GPR)

Gaussian processes (GPs) are used in machine learning as a nonparametric regression method for estimating continuous functions from sparse and noisy data [33]. For a GPIS, a training set consists of a set of input spatial locations $\mathcal{X} = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathbb{R}^d$, and signed distance observations $\mathbf{y} = \{y_1, \ldots, y_n\}, y_i \in \mathbb{R}$. In this work we will use $d = 2$ with input spatial locations $\mathcal{X}$ restricted to an $M \times M$ 2-dimensional grid with square cells. In practice, $\mathbf{y}$ might be acquired using KinectFusion, which forms an SDF from point clouds taken with an RGBD sensor [10], [30].

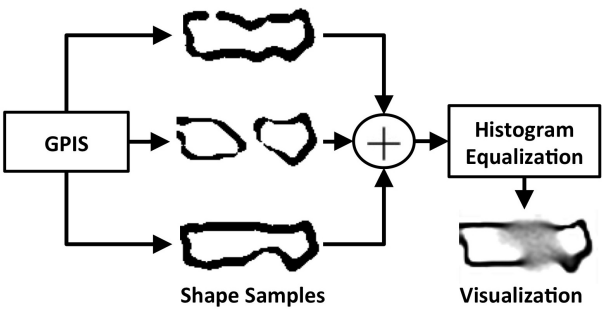

Fig. 2. Visualization method for 2D GPIS of Object H from Fig. 1 with shape uncertainty near the object center. (Left to right) We sample 1000 SDFs, threshold for zero crossings, average sample the samples together, and finally enhance contrast for easier perception of the uncertain regions.

A zero-mean Gaussian process implicit surface is specified by a covariance function $k(\cdot,\cdot)$, also referred to as a kernel, which measures the similarity in signed distance between spatial locations. Given a set of training data $\mathcal{D} = \{\mathcal{X}, \mathbf{y}\}$ and kernel $k(\cdot,\cdot)$, the posterior distribution on SDF f_* at a test location $\mathbf{x}_*$, $p(f_*|\mathbf{x}_*, \mathcal{D})$, is shown to be [33]:

$$p(f_* \mid \mathbf{x}_*, \mathcal{D}) \sim \mathcal{N}\big(\mu(\mathbf{x}_*), \sigma^2(\mathbf{x}_*)\big)$$
$$\mu(\mathbf{x}_*) = k(\mathcal{X}, \mathbf{x}_*)^\intercal (K + \sigma^2 I)^{-1}\mathbf{y} \quad (1)$$
$$\sigma^2(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathcal{X}, \mathbf{x}_*)^\intercal (K + \sigma^2 I)^{-1} k(\mathcal{X}, \mathbf{x}_*) \quad (2)$$

where $K \in \mathbb{R}^{n \times n}$ is a matrix with entries $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ and $k(\mathcal{X}, \mathbf{x}_*) = [k(\mathbf{x}_1, \mathbf{x}_*), \ldots, k(\mathbf{x}_n, \mathbf{x}_*)]^\intercal$. This derivation can also be used to predict the mean and variance of the SDF gradient by differentiating the kernel function, which can be used to obtain GPIS surface normals [11], [33], [37].

In this work we use the squared exponential kernel, as used by Dragiev et al.[11]. This kernel depends on two hyperparameters that can be set using maximum-likelihood estimation [41]. Other common kernels relevant to GPIS are the thin-plate splines kernel [41] and the Matern kernel [3].

B. Visualization

GPIS models can be visualized by sampling from the distribution on SDF zero-crossings, as illustrated in Fig. 2. We first sample many SDFs from the GPIS over our fixed spatial grid of input points, $\mathcal{X}$, and find the zero crossing of each sample SDF by searching for neighboring grid points with opposite signs. This results in a single shape 'image' on the spatial grid. We then average all sampled shapes and enhance the contrast using histogram equalization to make the uncertain areas easier to perceive [38]. The result is that regions of the shape that are more certain appear dark and crisp, while regions of high uncertainty appear blurry.

IV. Problem Statement

We consider grasping an object from above using a parallel-jaw gripper, using a 2-dimensional GPIS representation of the object contours to select grasps. We assume a known estimate of the coefficient of friction γ between the grasped object and the grippers, as well as hard contacts for the parallel jaws of the grippers. We also assume the object is rigid and stationary during the grasp.

A. Object Model

We assume a 2D GPIS representation that is augmented to predict the gradient of the SDF [37], as described in Section III. We use f to denote a SDF sample of the GPIS over the spatial grid, $f(\cdot) \sim \mathcal{N}(\mu(\cdot), \sigma^2(\cdot))$. We refer to the outward pointing surface normals at a spatial location $\mathbf{x} \in \mathbb{R}^2$ as $\mathbf{n}(\mathbf{x}) = \frac{\nabla\mu(\mathbf{x})}{\|\nabla\mu(\mathbf{x})\|_2}$, which is the normalized gradient of f. We assume that f will be evaluated over a fixed set of points on a grid, $\mathcal{X}$, as described in Section III. We also assume a uniform mass density of the object which can be used to predict the object center of mass $\mathbf{z} \in \mathbb{R}^2$ given the GPIS distribution.

B. Candidate Grasp Model

Our candidate grasp model is illustrated in Fig. 3. We consider desired grasps consisting of a spatial location for each of the parallel jaws relative to the GPIS representation of an object, $\mathbf{g}_1, \mathbf{g}_2 \in \mathbb{R}^2$, which also define the orientation of the grasp for parallel-jaw grippers. The gripper has a max opening of $w_g \in \mathbb{R}$ and that each jaw is $w_j \in \mathbb{R}$ wide. Therefore our set of *candidate grasps* is $\mathcal{G} = \{g = (\mathbf{g}_1, \mathbf{g}_2) : \|\mathbf{g}_1 - \mathbf{g}_2\| \leq w_g\}$, and we will henceforth refer to a single desired grasp as g.

In practice a robot may not be able to execute a desired grasp g exactly due to errors in trajectory following or registration to the object [18]. To handle this uncertainty, we define an *actual grasp* $\hat{g} = (\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2)$ as the location of the grasp during execution. We model the error in positioning the gripper as zero-mean Gaussian noise with isotropic covariance $\sigma_g^2 I$ about the desired location for the first jaw $\hat{\mathbf{g}}_1 \sim \mathcal{N}(\mathbf{g}_1, \sigma_g^2 I)$, with the location of the second jaw being conditionally dependent on the first. In practice σ_g^2 might be set based on repeatability measurements for a robot [28].

Given an actual grasp, we also define a *contact point* as the point at which the grasp comes into contact with an object when following the line segment between the two parallel jaws. We will refer to the contact configuration for an actual grasp $\hat{g}$ as $c = (\mathbf{c}_1, \mathbf{c}_2)$ where $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{R}^2$. Formally, a contact point $\mathbf{c}$ is the first 0-level set of an SDF f that the parallel jaws pass over when following approach direction $\mathbf{v} = \hat{\mathbf{g}}_2 - \hat{\mathbf{g}}_1$. This zero crossing satisfies $|f(\mathbf{c})| < \epsilon$ for some small $\epsilon \in \mathbb{R}, \epsilon > 0$. In practice the contacts can be found by testing points for the zero crossing condition along $\mathbf{v}$ [30]. We will also refer to the surface normals at contact configuration c as $n = (\mathbf{n}_1, \mathbf{n}_2)$ where $\mathbf{n}_1, \mathbf{n}_2 \in \mathbb{R}^2$.

C. Quality Measure

We measure the quality of a grasp using the probability of force closure P_F [20], [21], [23], [40]. To evaluate the P_F, we use the Ferrari-Canny grasp metric Q_F [12], [25], [27]. Given a set of contact wrenches $\mathcal{W} \in \mathbb{R}^6$ derived from contact locations c, normals n, and center of mass $\mathbf{z}$, $Q_F(c, n, \mathbf{z})$ measures the size of the largest ball around the origin in wrench space within the convex hull of $\mathcal{W}$ [12]. The force closure condition is equivalent to the positivity of Q_F, and therefore we define $P_F = P(Q_F > 0|g, \mu, \sigma^2)$ [40]. We will use henceforth use the notation $Q_F(\hat{g}, f)$ as shorthand

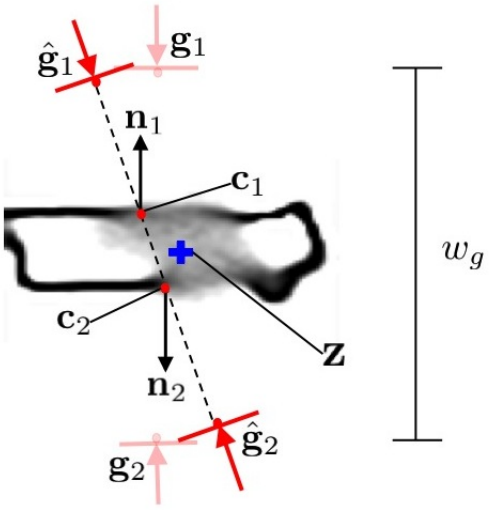

Fig. 3. Illustration of the variables defined for GP-GPIS-OPT on a Object H of Fig. 1 with shape uncertainty near the object center. Jaw placements are illustrated by a direction arrow and jaw. The target grasp $(\mathbf{g}_1, \mathbf{g}_2)$ of width w_g results in actual grasp $(\hat{\mathbf{g}}_1, \hat{\mathbf{g}}_2)$ due to errors in jaw placement. When the jaws close, the actual grasp contacts the surface at locations $(\mathbf{c}_1, \mathbf{c}_2)$ with outward-pointing surface normals $(\mathbf{n}_1, \mathbf{n}_2)$. Together with the center of mass $\mathbf{z}$, these values are sufficient to determine force closure for the grasp.

because the c, n, and $\mathbf{z}$ can be derived deterministically from actual grasp $\hat{g}$ and SDF sample f.

D. Grasp Planning Definition

The optimal grasp according to the P_F metric given a GPIS mean and variance function $\mu(\cdot), \sigma^2(\cdot)$ is:

$$g^* = \max_{g \in \mathcal{G}} P(Q_F(\hat{g}, f) > 0 \mid g, \mu, \sigma^2) \ . \tag{3}$$

We can evaluate this probability by taking the expectation of the indicator $\mathbb{1}(Q_F(\hat{g}, f) > 0)$ with respect to the distributions on actual grasps $\hat{g}$ and SDFs f:

$$\begin{aligned} P_F(g, \mu, \sigma^2) &= E[\mathbb{1}(Q_F(\hat{g}, f) > 0) \mid g, \mu, \sigma^2)] \\ &= \int_{\hat{\mathcal{G}}, \mathcal{F}} \mathbb{1}(Q_F(\hat{g}, f) > 0) p(\hat{g} \mid g) p(f \mid \mu, \sigma^2). \end{aligned}$$

where $\hat{\mathcal{G}}$ and $\mathcal{F}$ denote the sets of actual grasps and SDFs, respectively.

V. Grasp Optimization Algorithm

Our algorithm, GP-GPIS-OPT (Grasp Planner using GPIS OPTimization), finds locally optimal solutions of an approximation to Equation (3) using gradient-based optimization methods with finite differences to compute derivatives [32].

A. Quality Approximation

Direct optimization of Equation (3) requires a large number of shape samples to evaluate the objective for each grasp, which can be slow. However, we argue that under certain conditions on the SDF uncertainty at a contact, the probability of force closure on the the mean SDF μ, $\tilde{P}_F(g, \mu)$, may be a reasonable approximation of $P_F(g, \mu, \sigma^2)$.

Consider a desired grasp g and GPIS mean μ and variance σ^2 as defined in Section IV. Let $\tilde{c} = (\tilde{\mathbf{c}}_1, \tilde{\mathbf{c}}_2)$ denote the contact points for a grasp g on the mean SDF μ. Let $A_1(f)$ denote the event that the parallel jaws pass over a zero crossing on SDF sample f before reaching $\tilde{\mathbf{c}}_1$ when attempting g. Also, let $B_1(f)$ denote the event that $\tilde{\mathbf{c}}_1$ is not a zero crossing of SDF sample f, which is defined by $|f(\tilde{\mathbf{c}}_1)| > \epsilon$, as described in Section IV-B.

Now suppose that $P(A_1(f)) < \delta$ for some $\delta \in [0, 1]$. Also suppose that $\sigma^2(\tilde{\mathbf{c}}_1) < \tau^2$ for some $\tau \in \mathbb{R}$ and $\tau << \epsilon$. Then for any SDF f we have $|f(\mathbf{c}_1)| \leq |\mu(\tilde{\mathbf{c}}_1) \pm 2\sigma(\tilde{\mathbf{c}}_1)| < \epsilon + 2\tau \approx \epsilon$ with approximately 95% probability, since $\mathcal{N}(\mu(\tilde{\mathbf{c}}_1), \sigma^2(\tilde{\mathbf{c}}_1))$ is a 1-dimensional Gaussian [29]. Therefore $\tilde{\mathbf{c}}_1$ is still likely to be a zero crossing for other sample SDFs from the GPIS. Furthermore, the probability that we do not contact the surface at $\tilde{\mathbf{c}}_1$ when attempting grasp g is:

$$\begin{aligned} P(A_1(f) \cup B_1(f)) &\leq P(A_1(f)) + P(B_1(f)) \\ &\leq \delta + (1.0 - 0.95) \\ &\leq \delta + 0.05 \end{aligned}$$

Thus, given our assumptions $\tilde{\mathbf{c}}_1$ is likely to be the same contact for approximately $(0.95 - \delta)$ of SDFs sampled from the GPIS. Finally, assuming that $\mathbf{z} \approx E[\mathbf{z}|\mu, \sigma^2]$, $\tilde{P}_F(g, \mu)$ may be a reasonable approximation of $P_F(g, \mu, \sigma^2)$.

We attempt to satisfy the assumption $\sigma^2(\tilde{\mathbf{c}}_1) < \tau^2$ by using a soft version of the constraint $\sigma^2(\tilde{\mathbf{c}}_i) < \tau^2$ in our optimization. Specifically, we propose to penalize the σ^2 by $\lambda \in \mathbb{R}, \lambda > 0$ which has one-to-one correspondence with a choice of τ [4]:

$$\underset{g \in \mathcal{G}}{\text{maximize}} \ \tilde{P}_F(g, \mu) - \lambda(\sigma^2(\tilde{\mathbf{c}}_1) + \sigma^2(\tilde{\mathbf{c}}_2)) \tag{4}$$

High values of λ, which correspond to smaller τ, may increase the accuracy of the approximation but may discard grasps with high P_F. Very small values of λ may increase the set of possible grasps but the approximation will become increasingly inaccurate. We acknowledge that the choice of λ is somewhat arbitrary and in practice λ might be set using cross-validation, similar to choosing a regularization penalty in regression models [4], [15].

We cannot directly incorporate minimization of the δ satisfying $P(A_1(f)) < \delta$ at this time and consider it a current shortcoming of our method. Therefore, this approximation may be poor is when there is considerable uncertainty in the space between the locations of a desired grasp $\mathbf{g}_1, \mathbf{g}_2$ and the zero crossing of μ. However, in Section VI-B we show that empirically this does not prevent GP-GPIS-OPT from finding grasps with high P_F.

B. Grasp Constraints

Parallel-jaw grippers have two points-of-contact which must be approached from opposite directions, and thus we constrain the grasp contacts on the mean shape to be antipodal [5]. For unit outward pointing normal vectors n at contact c and approach direction $\mathbf{v} = \hat{\mathbf{g}}_1 - \hat{\mathbf{g}}_2$, c is an antipodal pair if the contacts have (a) opposite normals and

(b) normals aligned with the approach direction [5]. We convert these constraints to inequalities because on actual objects it may be impossible to satisfy the definitions exactly:

$$\|\mathbf{n}_1 + \mathbf{n}_2\|_2^2 \leq \alpha \quad (5)$$

$$\mathbf{n}_1^T \mathbf{v} \geq \beta \|\mathbf{v}\|_2^2 \quad (6)$$

$$-\mathbf{n}_2^T \mathbf{v} \geq \beta \|\mathbf{v}\|_2^2 \quad (7)$$

For 2-dimensional shapes the choice of $\beta = \cos(\arctan(\gamma))$ guarantees that points in the feasible region will be in force closure for the mean SDF [5], [7], and α may be set using a grid search.

C. Grasp Selection Algorithm

Taking the objective of Equation 4 and the constraints of Equations 5, 6, and 7 yields the optimization objective:

$$\begin{aligned} \underset{g \in \mathcal{G}}{\text{maximize}} \quad & \tilde{P}_F(g, f) - \lambda(\sigma^2(\mathbf{c}_1) + \sigma^2(\mathbf{c}_2)) \\ \text{subject to:} \quad & \|\mathbf{n}_1 + \mathbf{n}_2\|_2^2 < \alpha \\ & \mathbf{n}_1^T \mathbf{v} > \beta \|\mathbf{v}\|_2^2 \\ & -\mathbf{n}_2^T \mathbf{v} > \beta \|\mathbf{v}\|_2^2 \end{aligned}$$

where $\mathbf{c}_1, \mathbf{c}_2, \mathbf{n}_1, \mathbf{n}_2$, and $\mathbf{v}$ are derived from candidate grasp g on the mean shape μ as described in Section IV-B.

This problem is non-convex due to the Q_F quality evaluation, and therefore we can only expect a locally optimal solution to this problem for a single random initialization. The GP-GPIS-OPT (Grasp Planner using GPIS OPTimization) algorithm, detailed in Algorithm 1, repeatedly finds locally optimal grasps for Equation (8) from a user-specified number of random initializations N_g and selects the grasp with the highest objective value. In this work we use $N_g = 20$, based on the empirical worst-case time to converge to a solution for 25×25 GPIS models.

GP-GPIS-OPT uses Sequential Convex Programming (SCP) to find a locally optimal grasp for Equation (8) given a random initial grasp. SCP iteratively forms a convex approximation to the problem and solves the approximation within a confidence [36], [42]. SCP turns non-convex constraints into penalties and iteratively increasing a penalty coefficient, similar to Interior Point methods [36]. Although P_F is not differentiable everywhere, we use finite differences to approximate the gradients, as has been shown to work well empirically for optimizing Q_F on deterministic shapes [32].

VI. Experiments

A. Dataset

We evaluated the performance of our grasp selection method on a set of 8 household objects. Our test set of objects is displayed in Fig. 1 and is available at `http://rll.berkeley.edu/grasping/`. The objects chosen illustrate various properties which lead to missing or invalid measurements with an RGBD camera: (a) transparency, (b) translucency, (c) specularity, and (d) partial visibility. Missing measurements arise from specularity or partial visibility, and appear as black regions in point clouds. For example, the wings of Object B and the metal parts of

Input: GPIS Model $\mu(\cdot), \sigma^2(\cdot)$, Grid $\mathcal{X}$, and Number of Initial Grasps N_g
Result: Grasp Proposal g^* and Objective Value V^*
Initialize mean shape $\mu(\mathcal{X})$;
Initialize grasp count $k = 0$;
Initialize $V^* = -\infty$;
while $k < N_g$ **do**
  Sample valid initial grasp $g_{0,k}$;
  Use SCP to find a locally optimal grasp g_k for Equation 8 with objective value V_k ;
  if $V_k > V^*$ **then**
    $g^* = g_k$;
    $V^* = V_k$;
  end
  k = k + 1 ;
end

Algorithm 1: The GP-GRASP-OPT Algorithm

Object C and G cannot be sensed by the RGBD camera. Invalid measurements, such as measuring the table behind a transparent region, also occur for the Objects A and E.

We created a 2-dimensional SDF for each object in a 25×25 grid based on the point clouds of the object by specifying the shape surface as the SDF zero-crossing. We also created an occupancy map, which holds 1 if the point could be observed with a depth sensor and 0 if the point could not be observed, and a measurement noise map, which holds the variance of 0-mean noise added to the SDF values. The noise values were set uniformly for observed locations. The parameters of the squared exponential kernel were selected using maximum-likelihood on a held-out set of validation shapes; see [33] for details on this method.

B. Grasp Planning

To evaluate the performance of GP-GPIS-OPT, we compared the probability of force closure P_F performance of 5 competing methods for selecting a grasp g:

- **GP-M:** Selects g with the highest Ferrari-Canny quality Q_F on the mean SDF from 1000 random grasp samples.
- **GP-P:** Selects g with the highest P_F from 1000 randomly sampled grasps, evaluating each grasp with Monte-Carlo integration over 1000 shapes [7], [21].
- **GP-D:** Optimizes P_F on the mean SDF μ.
- **GP-U:** Optimizes P_F on the mean SDF μ with uncertainty penalty λ, as in Equation (4).
- **GP-G:** The GP-GPIS-OPT algorithm, as detailed in Algorithm 1.

GP-M is fast because it ignores shape uncertainty, but may find grasps with low P_F. GP-P selects grasps with high P_F but is slow because it requires a large number of SDF samples and grasps [7], [20], [21]. This method may be prohibitively slow on higher resolution models; for example, sampling n shapes from a fixed grid of size $m \times m$ takes $O(nm^2)$ space to store samples and $O(m^6)$ time because it involves the Cholesky decomposition of an m^2-dimensional Gaussian [24], [29]. GP-GPIS-OPT is designed to be faster than GP-P and is also to select grasps with high P_F. GP-D and GP-U are of comparable speed and are presented

	Probability of Force Closure P_F						Runtime (sec)	
Object	GP-M	GP-I	GP-D	GP-U	GP-P	**GP-G**	GP-P	**GP-G**
A	0.72	**0.99**	0.43	0.13	**0.99**	**0.99**	645.5	**110.9**
B	0.38	0.56	0.12	0.57	**0.98**	0.79	715.8	**72.4**
C	0.15	0.06	0.32	0.50	**0.96**	0.90	744.7	**85.5**
D	0.75	0.07	0.00	0.01	0.91	0.37	848.7	**84.3**
E	0.47	0.98	0.20	**0.99**	**0.99**	**0.99**	545.9	**149.8**
F	0.70	0.46	0.00	0.98	**0.99**	**0.99**	701.8	**75.6**
G	0.98	0.77	0.83	0.00	**0.99**	**0.99**	817.6	**107.5**
H	0.42	0.62	0.47	0.00	**0.99**	0.93	715.0	**85.5**

TABLE I

COMPARISON OF THE PROBABILITY OF FORCE CLOSURE P_F (A) GP-M, WHICH SELECTS A GRASP BASED ON ONLY THE MEAN SDF, (B) GP-I, WHICH SELECTS THE INITIAL GRASP IN THE OPTIMIZATION WITH HIGHEST P_F, (C) GP-D, WHICH OPTIMIZES P_F ON THE MEAN SDF, (D) GP-U, WHICH OPTIMIZES THE P_F WITH AN UNCERTAINTY PENALTY, AND (E) GP-G, THE GP-GPIS-OPT ALGORITHM. WE ALSO COMPARE WITH GP-P, WHICH CHOOSES THE GRASP WITH HIGHEST P_F FROM A SET OF 1000 RANDOM GRASP SAMPLES. WE ASSUME GP-P TO BE NEAR-GROUND-TRUTH DUE TO THE LARGE NUMBER OF GRASP SAMPLES. GP-GPIS-OPT PERFORMS AS WELL AS GP-P ON 4 OF 8 OBJECTS AND SELECTS GRASPS WITH HIGHER P_F THAN METHODS OTHER THAN GP-P FOR 7 OF 8 OBJECTS. WE ALSO COMPARE THE WORST-CASE RUNTIME TO CONVERGE TO A GRASP PLAN FOR GP-GPIS-OPT AND GP-P, AND SEE GP-GRASP-OPT IS ON AVERAGE 7.9× FASTER THAN GP-P.

to evaluate the use of the uncertainty penalty and antipodal constraints in GP-GPIS-OPT.

The parameters of GP-GRASP-OPT, GP-D, and GP-U were $N_g = 20$ and a grasp approach uncertainty of $\sigma_g^2 = 0.25$. The uncertainty penalty of GP-GRASP-OPT and GP-U was set to $\lambda = 2.0$ based on a grid search using the set of validation shapes [15]. The sampling-based methods GP-M and GP-P used 1000 random grasp samples, chosen based on the number of random samples required for the methods to converge to a single grasp. All experiments were run on machine with OS X with a 2.7 GHz Intel core i7 processor and 16 GB 1600 MHz memory in Matlab 2014a.

Table I compares P_F for each of the grasp selection methods and the runtimes for GP-P and GP-GPIS-OPT. To illustrate progress made by the optimization we also include the grasp chosen by GP-I, the random initial grasp used in GP-GPIS-OPT with the highest P_F. Since GP-P samples an exhaustive set of 1000 grasps, we consider it near-optimal for the purposes of comparing P_F. The runtimes reported for GP-P and GP-G are the worst-case runtimes to converge to a grasp plan. For GP-P this is the time to evaluate 753 grasp samples, the worst case number observed over all experiments, and for GP-G this is the time to optimize all N_g grasps.

We see that GP-GPIS-OPT chooses the grasp with the same P_F as GP-P on 4 of 8 objects and the grasp with the highest P_F for methods other than GP-P on 7 of 8 objects. GP-GRASP-OPT is also faster than GP-P on all objects, with approximately a 7.9× speedup over GP-P on average.

Fig. 4 illustrates grasps chosen by GP-M, GP-P, and GP-GPIS-OPT for Objects A-D and compares P_F and Q_N, the Ferrari-Canny quality on the nominal shape. We see that for several objects, GP-M may choose a grasp that is not in force closure on the nominal (true) shape. For example, on Object A the grasp chosen by GP-M is not in force closure because it ignores the tape dispenser, which cannot be sensed with the depth sensor. However, both grasps that take into account shape uncertainty are able to avoid the dispenser and choose a grasp with high P_F and Q_N. We also see that grasps chosen by our algorithm have higher Q_N than GP-M on all shapes. Additionally, GP-GPIS-OPT chooses grasps qualitatively and quantitatively similar to GP-P but is faster, as reported in Table I.

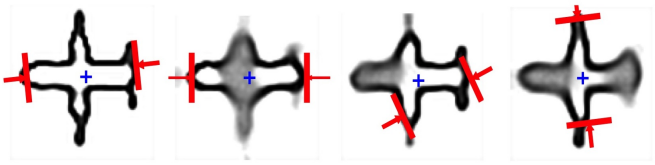

Fig. 5. Comparison of grasps chosen by our algorithm on Object B, the toy plane, with different patterns of shape uncertainty: (left to right) uncertainty in the tips and edges of wings, uncertainty in the entire wings, uncertainty in the cockpit end of the plane, and uncertainty in both the cockpit and rudder end of the plane.

C. Sensitivity to Shape Uncertainty

Fig. 5 compares grasps chosen by GP-GRASP-OPT with different patterns of shape uncertainty on Object B, the toy plane. We consider shape uncertainty when Object B, the toy plane, has noise only on the edges of the wings, when the entire wings are not observed, when the tip of the plane is not observed, and when both the front and the back are not observed. In both cases when there is uncertainty in the wings, the selected grasp is on the endpoints of the plane. With uncertainty on the front end of the plane the algorithm chooses a grasp that leverages a wing and a rear stabilizer, but with uncertainty on both endpoints the algorithm selects the tip of the wings, resulting in an unstable grasp.

D. Physical Grasp Experiments with the PR2

We tested grasps for Object A, a roll of Scotch tape, on the two-armed Willow Garage PR2 [1]. We created a GPIS from a point cloud segmentation of the object. We compared GP-M, GP-P, and GP-GPIS-OPT with $N_g = 10$ and $\lambda = 2.0$, chosen by grid search [15]. The grasps chosen by these methods are illustrated in Fig. 6. The grasp chosen by GP-M had predicted $P_F = 0.48$ and achieved force closure on

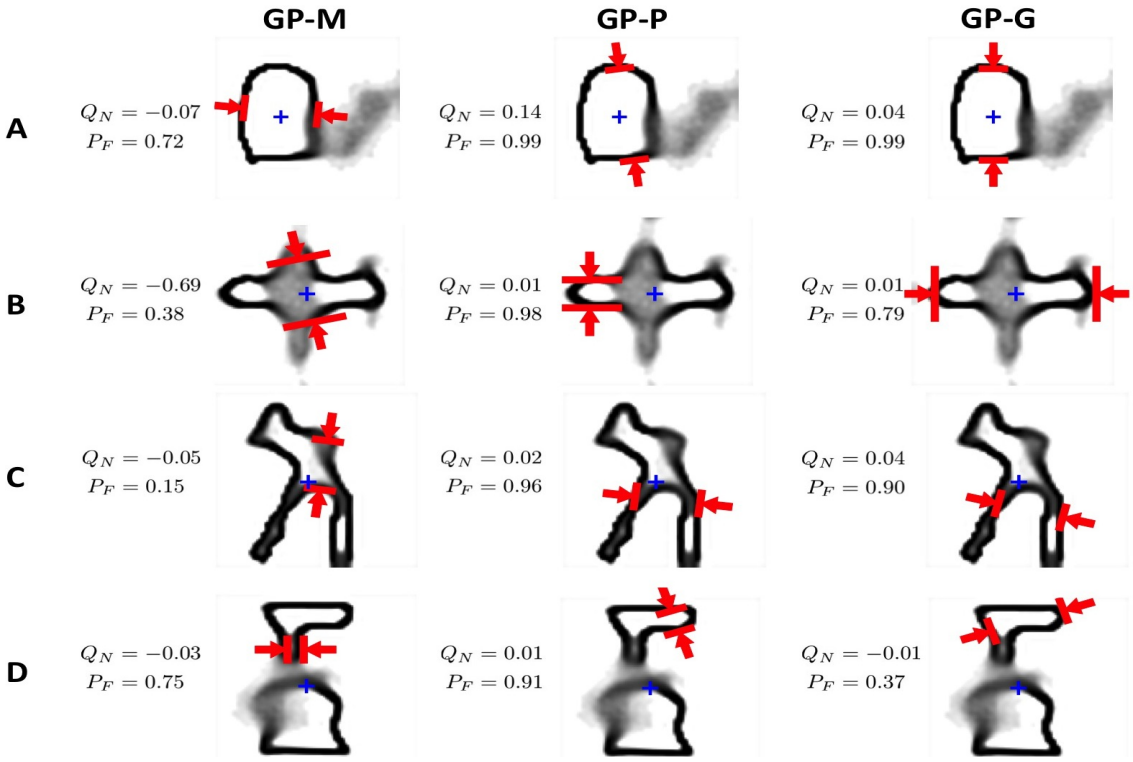

Fig. 4. Visual comparison of grasps on Objects A, B, C, and D (top to bottom) selected by (left to right) (i) GP-M, which chooses the grasp with highest Q_F on the mean SDF, (ii) GP-P, which chooses the grasp with the highest probability of force closure P_F from 1000 random grasp samples and (iii) GP-GPIS-OPT, our algorithm for selecting grasps using optimization. The Ferrari-Canny quality on the nominal shape Q_N and P_F are listed to the left of each grasp. We see that GP-GPIS-OPT outperforms GP-M in terms of Q_N; for example, on the Object A GP-M chooses a grasp on the transparent tape dispenser but GP-GPIS-OPT avoids this region due to shape uncertainty. Also, grasps chosen by GP-GRASP-OPT are comparable to those chosen by GP-P in both P_F and Q_N, but GP-GPIS-OPT requires fewer random initializations.

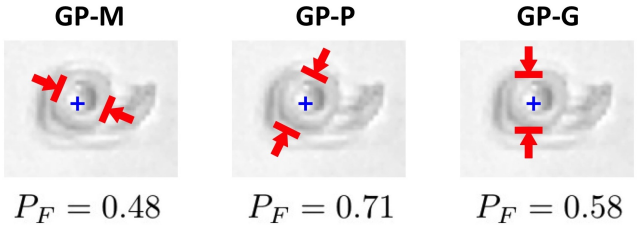

Fig. 6. Grasps selected for execution on the PR2 for Object A overlayed on images of the object from a head-mounted Primesense Carmine. The grasp with highest Q_F on the mean shape fails does not consider the transparent dispenser, but GP-GPIS-OPT avoids the area due to uncertainty in the object geometry near the dispenser.

only 1 of 10 trials. The failures were due to the gripper contacting and sliding along the transparent outer surface of the tape canister, which could not be sensed reliably with either depth or color images. The grasp chosen by GP-P had predicted $P_F = 0.71$ and achieved force closure on 10 of 10 trials. The grasp chosen by GP-GPIS-OPT had $P_F = 0.58$ and also achieved force closure on 10 of 10 trials.

VII. Discussion and Future Work

We presented GP-GPIS-OPT, an algorithm for selecting parallel-jaw grasps with high probability of force closure P_F on Gaussian process implicit surface representations of shape uncertainty. Our experiments suggest that GP-GPIS-OPT plans grasps with higher P_F than methods that ignore shape uncertainty, and it is faster than the most common existing method for grasp planning with shape uncertainty.

Future work will extend our method to 3D GPIS models, which will bring several challenges. Since GPIS construction takes $O(n^3)$ time where n is the number of input spatial locations, we will investigate how to construct GPIS models efficiently by using a representative subset of training points [16]. We will also implement GP-GPIS-OPT on a multicore or Cloud Computing framework and assess performance on 3D models, as it is parallelizable over random initial grasps (Line 5 of Algorithm 1). Furthermore, we will extend our method to multi-point grasps on 3D GPIS models by optimizing over joint angles and using forward kinematics to determine contact locations [9].

A current shortcoming of GP-GPIS-OPT is the use of a heuristic approximation to P_F that introduces a penalty term λ into our objective. Future work will study the effect of λ on the quality of selected grasps and consider alternative quality measures and approximations. We will also consider re-ranking grasps optimized by GP-GPIS-OPT by the exact quality P_F using Multi-Armed Bandit sampling methods [24].

Our dataset, code, and videos of our experiments are available at `http://rll.berkeley.edu/grasping/`.

VIII. Acknowledgments

This work has been supported in part by the U.S. National Science Foundation under Award IIS-1227536, and by grants from Google and Cisco. We thank our colleagues who gave feedback, suggestions, and help on the PR2, in particular James Kuffner, Florian Pokorny, Michael Laskey, Greg Kahn, Sanjay Krishnan, Sylvia Herbert, Zoe McCarthy, Alex Turney, John Schulman, and Dylan Hadfield-Menell.

References

[1] "Willow garage pr2," http://www.willowgarage.com/pages/pr2/overview.

[2] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *Robotics, IEEE Transactions on*, vol. 27, no. 3, pp. 616–629, 2011.

[3] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 3180–3186.

[4] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2009.

[5] I.-M. Chen and J. W. Burdick, "Finding antipodal point grasps on irregularly shaped objects," *Robotics and Automation, IEEE Transactions on*, vol. 9, no. 4, pp. 507–512, 1993.

[6] J.-S. Cheong, H. Kruger, and A. F. van der Stappen, "Output-sensitive computation of force-closure grasps of a semi-algebraic object," *Automation Science and Engineering, IEEE Transactions on*, vol. 8, no. 3, pp. 495–505, 2011.

[7] V. N. Christopoulos and P. Schrater, "Handling shape and contact location uncertainty in grasping two-dimensional planar objects," in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*. IEEE, 2007, pp. 1557–1563.

[8] V. N. Christopoulos and P. R. Schrater, "Grasping objects with environmentally induced position uncertainty," *PLoS computational biology*, vol. 5, no. 10, p. e1000538, 2009.

[9] M. Ciocarlie, C. Goldfeder, and P. K. Allen, "Dimensionality reduction for hand-independent dexterous robotic grasping," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems: San Diego, CA, 29 October-2 November 2007*. IEEE, 2007, pp. 3270–3275.

[10] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM, 1996, pp. 303–312.

[11] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2011, pp. 2845–2850.

[12] C. Ferrari and J. Canny, "Planning optimal grasps," in *Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on*. IEEE, 1992, pp. 2290–2295.

[13] K. Y. Goldberg and M. T. Mason, "Bayesian grasping," in *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*. IEEE, 1990, pp. 1264–1269.

[14] K. Y. Goldberg, "Stochastic plans for robotic manipulation," 1990.

[15] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.

[16] A. Gotovos, N. Casati, G. Hitz, and A. Krause, "Active learning for level set estimation," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

[17] G. A. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme, "Uncertainty-driven view planning for underwater inspection," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4884–4891.

[18] K. Hsiao, M. Ciocarlie, and P. Brook, "Bayesian grasp planning," in *ICRA 2011 Workshop on Mobile Manipulation: Integrating Perception and Manipulation*, 2011.

[19] J. Ilonen, J. Bohg, and V. Kyrki, "Fusing visual and tactile sensing for 3-d object reconstruction while grasping," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 3547–3554.

[20] B. Kehoe, D. Berenson, and K. Goldberg, "Estimating part tolerance bounds based on adaptive cloud-based grasp planning with slip," in *Automation Science and Engineering (CASE), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1106–1113.

[21] ——, "Toward cloud-based grasping with uncertainty in shape: Estimating lower bounds on achieving force closure with zero-slip push grasps," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 576–583.

[22] J. Kim, K. Iwamoto, J. J. Kuffner, Y. Ota, and N. S. Pollard, "Physically-based grasp quality evaluation under uncertainty," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3258–3263.

[23] J. Laaksonen, E. Nikandrova, and V. Kyrki, "Probabilistic sensor-based grasping," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 2019–2026.

[24] M. Laskey, Z. McCarthy, J. Mahler, F. Pokorny, S. Patil, J. van den Berg, D. Kragic, P. Abbeel, and K. Goldberg, "Multi-armed bandit models for sample-based grasp planning in the presence of uncertainty." in *Robotics and Automation (ICRA), 2015 IEEE International Conference on (Under Review)*. IEEE, 2015.

[25] B. León, S. Ulbrich, R. Diankov, G. Puche, M. Przybylski, A. Morales, T. Asfour, S. Moisio, J. Bohg, J. Kuffner, *et al.*, "Opengrasp: a toolkit for robot grasping simulation," in *Simulation, Modeling, and Programming for Autonomous Robots*. Springer, 2010, pp. 109–120.

[26] G. Liu, J. Xu, X. Wang, and Z. Li, "On quality functions for grasp synthesis, fixture planning, and coordinated manipulation," *Automation Science and Engineering, IEEE Transactions on*, vol. 1, no. 2, pp. 146–162, 2004.

[27] A. T. Miller and P. K. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics & Automation Magazine, IEEE*, vol. 11, no. 4, pp. 110–122, 2004.

[28] B. Mooring and T. Pack, "Determination and specification of robot repeatability," in *Robotics and Automation. Proceedings. 1986 IEEE International Conference on*, vol. 3. IEEE, 1986, pp. 1017–1023.

[29] D. F. Morrison, "Multivariate statistical methods. 3," *New York, NY. Mc*, 1990.

[30] R. A. Newcombe, A. J. Davison, S. Izadi, P. Kohli, O. Hilliges, J. Shotton, D. Molyneaux, S. Hodges, D. Kim, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on*. IEEE, 2011, pp. 127–136.

[31] F. Panahi, M. Davoodi, and A. F. van der Stappen, "Orienting parts with shape variation," 2014.

[32] F. T. Pokorny, K. Hang, and D. Kragic, "Grasp moduli spaces." in *Robotics: Science and Systems*, 2013.

[33] C. E. Rasmussen, "Gaussian processes for machine learning," 2006.

[34] M. A. Roa and R. Suárez, "Computation of independent contact regions for grasping 3-d objects," *Robotics, IEEE Transactions on*, vol. 25, no. 4, pp. 839–850, 2009.

[35] ——, "Finding locally optimum force-closure grasps," *Robotics and Computer-Integrated Manufacturing*, vol. 25, no. 3, pp. 536–544, 2009.

[36] J. Schulman, J. Ho, A. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization." in *Robotics: Science and Systems*, vol. 9, no. 1. Citeseer, 2013, pp. 1–10.

[37] E. Solak, R. Murray-Smith, W. E. Leithead, D. J. Leith, and C. E. Rasmussen, "Derivative observations in gaussian process models of dynamic systems," 2003.

[38] J. A. Stark, "Adaptive image contrast enhancement using generalizations of histogram equalization," *Image Processing, IEEE Transactions on*, vol. 9, no. 5, pp. 889–896, 2000.

[39] F. Steinke, B. Schölkopf, and V. Blanz, "Support vector machines for 3d shape processing," in *Computer Graphics Forum*, vol. 24, no. 3. Wiley Online Library, 2005, pp. 285–294.

[40] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 557–562.

[41] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," *Gaussian Proc. in Practice*, 2007.

[42] S. Wright and J. Nocedal, *Numerical optimization*. Springer New York, 1999, vol. 2.