

Toward An Analytic Theory of Intrinsic Robustness for Dexterous Grasping

Albert H. Li[†], Preston Culbertson[‡], Aaron D. Ames^{†,‡}

Abstract—Conventional approaches to grasp planning require perfect knowledge of an object’s pose and geometry. Uncertainties in these quantities induce uncertainties in the quality of planned grasps, which can lead to failure. Classically, grasp robustness refers to the ability to resist external disturbances *after* grasping an object. In contrast, this work studies robustness to intrinsic sources of uncertainty like object pose or geometry affecting grasp planning *before* execution. To do so, we develop a novel analytic theory of grasping that reasons about this *intrinsic robustness* by characterizing the effect of friction cone uncertainty on a grasp’s force closure status. We apply this result in two ways. First, we analyze the theoretical guarantees on intrinsic robustness of two grasp metrics in the literature, the classical Ferrari-Canny metric and more recent min-weight metric. We validate these results with hardware trials that compare grasps synthesized with and without robustness guarantees, showing a clear improvement in success rates. Second, we use our theory to develop a novel analytic notion of *probabilistic* force closure, which we show can generate unique, uncertainty-aware grasps in simulation.

I. INTRODUCTION

Grasp synthesis has been a canonical problem in robotic manipulation since the field’s inception. Despite decades of work, dexterous grasping is still challenging, since multi-finger hands are high-dimensional and have complex kinematics. Broadly, two approaches toward dexterous grasping exist. *Analytic* methods maximize a grasp metric that measures grasp quality [1]. Such methods are principled and intuitive, but suffer from two main drawbacks: (a) they usually require perfect knowledge of an object, and (b) efficiently optimizing metrics is typically hard while enforcing kinematic and collision constraints [2]. In response, many *data-driven* methods for grasping account for uncertainty implicitly using empirical grasp data, but lack guarantees or interpretability, often checking constraint satisfaction *post hoc* rather than enforcing them during synthesis [3], [4]. Since high-quality, feasible grasps are rare to find randomly, these methods usually require many grasp data for training, which may be difficult to acquire (particularly from hardware).

Classically, grasp robustness reports the size of *extrinsic* disturbances a grasp can resist post-execution, like an unexpected force on the object [5]. But, many uncertainties we call *intrinsic* (like those in object pose or geometry) affect planning *before* execution. Poor plans create brittle grasps, so intrinsic uncertainties reduce extrinsic robustness. Despite




Fig. 1: This work develops an analytic theory of uncertainty-aware grasping, which enables an analysis of the uncertainty tolerance of classical grasp metrics, and the development of new, probabilistic approaches. Above we plot a grasp obtained by optimizing our proposed probabilistic metric, PONG, on a toy example with varying surface normal uncertainty. While there exist many robust grasps for the true geometry, the grasp synthesized using our probabilistic metric places the fingertips in the minimum-uncertainty region.

this, intrinsic robustness remains understudied, and the aim of this work is to provide new theory to analyze it.

Our main contribution is a mathematical sufficient condition that guarantees intrinsic grasp robustness under a linear friction cone model. This immediately proves that the classical Ferrari-Canny metric ε [5], designed to characterize resistance to extrinsic disturbances, also measures intrinsic robustness (Theorem 1 and Corollary 1). We further apply our result in two ways to demonstrate its utility.

First, we prove the recently proposed *min-weight metric* ℓ^* , computable via solving an LP, [6] uniformly lower bounds (a positive multiple of) the Ferrari-Canny metric ε , so ℓ^* grants intrinsic robustness and maximizing it is justified. We validate the theory by comparing ℓ^* -synthesized grasps against a competitive baseline on dozens of hardware trials, showing superior performance.

Second, we develop a novel *analytic* notion of probabilistic force closure, and develop an uncertainty-aware metric called *probabilistic object normals for grasping* (PONG). We demonstrate that maximizing PONG on objects with synthetic belief distributions leads to probabilistically-robust grasps (e.g., Fig. 1). In simulation, we show that ℓ^* performs similarly to PONG on uncertain objects, showing that the aforementioned bound is not conservative in practice.

A. Related Work

While this paper presents an analytic treatment of intrinsic uncertainty, prior work has studied it from data-driven perspectives. For example, [7] notes that ε is sensitive to uncertainty in object pose and estimates force closure probability via Monte Carlo simulation to rank grasps by robustness. *Gaussian process implicit surfaces* (GPISs) [8]

[†] A. H. Li and A. D. Ames are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA, {alberthli, ames}@caltech.edu.

[‡] P. Culbertson and A. D. Ames are with the Department of Civil and Mechanical Engineering, California Institute of Technology, Pasadena, CA 91125, USA, {pculbert, ames}@caltech.edu.

have represented uncertain geometry for parallel-jaw grasping [9], dexterous grasping with tactile sensors [10], and control synthesis [11]. But, GPSs are difficult to supervise from sensor data (requiring known signed-distance labels), scale poorly with data quantity, and lack expressivity (due to strong smoothness priors imposed by usual kernels).

An alternate approach is to directly learn a probabilistic metric from data. For example, Dex-Net 2.0 predicts the robustness of a batch of planar parallel-jaw grasps, selecting the most robust one [12]. Similar works have optimized multi-finger grasps with gradient-based optimization by leveraging the differentiability of neural networks [13], [14].

Many works have proposed differentiable approximations of analytic metrics or the force closure condition for gradient-based grasp synthesis. Such methods include solving a large sum of squares program [15]; optimizing a penalty-based relaxation of the force closure condition [16], [17]; solving a bilinear optimization program with a QP force closure constraint [18]; and maximizing a differentiable proxy for the Ferrari-Canny metric called the min-weight metric ℓ^* [6]. However, it is typically not understood how well or whether these approximations preserve grasp robustness.

We use our results to prove that maximizing ℓ^* , the fastest of the above methods, equivalently maximizes intrinsic robustness, which theoretically justifies its use as a fast and principled approximation of the Ferrari-Canny metric. We note that while other works have derived lower bounds for ε (e.g., [19], [20], [15]), the cost of computing these metrics is significantly higher than for ℓ^* , which requires only solving a small linear program. In future work, we hope to similarly analyze the robustness properties of other metrics.

B. Mathematical Preliminaries

We consider a fixed-base, fully-actuated rigid-body serial manipulator with a multi-finger hand and configuration $q \in \mathcal{Q}$. Assume the hand has n_f fingers contacting the object at points $\{x^i\}_{i=1}^{n_f}$ with inward pointing surface normals $\{n^i\}_{i=1}^{n_f}$ and forward kinematics maps $x^i = FK^i(q)$. Denote the compact manipuland $\mathcal{O} \subset \mathbb{R}^3$ with surface $\partial\mathcal{O}$.

As shorthand, let $\mathcal{C}_x := \text{conv}(\{x_l\})$ denote the convex hull of a finite set of points indexed by l . Define the wedge operator $[\cdot]_{\times} : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ such that $[a]_{\times} b := a \times b$ for $a, b \in \mathbb{R}^3$. We denote an optimized *feasible* configuration as q^* , meaning no undesired collisions while contacting the object. The fingers are modeled as point contacts with friction.

Recall that for Coulomb friction, a contact force f satisfies the *no-slip* condition if for friction coefficient μ , $\|f_t\| \leq \mu \cdot f_n$, where t and n denote tangent and (positive) normal components of f . We call such forces *Coulomb-compliant*. A force applied at point x induces the torque $\tau = x \times f$, which in turn induces a corresponding wrench $w = (f, \tau)$.

The *friction cone* at a point x is the cone centered about a surface normal n consisting of all Coulomb-compliant forces. We use a standard pyramidal approximation of it with n_s sides, calling its edges *basis forces*. We emphasize that the induced *basis wrenches* depend entirely on these pyramids via the basis forces, which depend on a normal n at a point x ;

so, we always implicitly have $w = w(q)$ and we refer to basis wrench uncertainty equivalently as *friction cone uncertainty*. Since it depends on x and n , the friction cone implicitly captures the effects of intrinsic uncertainty on grasps.

We let $n_w = n_f n_s$ denote the number of basis wrenches (indexed w_j^i), each associated with finger i and pyramid edge j . For brevity, we often combine indices i, j into a single one l . Let $\mathcal{I} = \{1, \dots, n_f\}$, $\mathcal{J} = \{1, \dots, n_s\}$, $\mathcal{L} = \{1, \dots, n_w\}$ denote index sets, and $\mathcal{W} = \{w_l\}_{l=1}^{n_w}$ denote a set of basis wrenches. Optimal solutions to optimization programs are denoted with asterisks. A set is *degenerate* if it has zero volume. $B_r(c)$ denotes a ball of radius r centered at c .

II. A THEORY OF INTRINSICALLY-ROBUST GRASPING

Recall that a grasp is force closure if it can generate Coulomb-compliant forces on the object to resist any disturbance wrench. A sufficient condition for force closure is that the origin is contained in the convex hull of a grasp's basis wrenches, i.e., $0 \in \mathcal{C}_w$ [6]. We say that such basis wrenches *certify* force closure.

Since the basis wrenches depend on contacts x and surface normals n , force closure is a function of both an object's geometry and a grasp's contacts. Thus, intrinsic uncertainties generate uncertainty in a grasp's basis wrenches, also reducing extrinsic robustness. In this section, we analyze how much uncertainty a nominally force closure can tolerate while still remaining force closure.

A. Characterizing a Grasp's Uncertainty Tolerance

Consider a nominal set of basis wrenches $\mathcal{C}_{\bar{w}}$ and the convex hull of a perturbed version, \mathcal{C}_w (see Fig. 2a/b). For a given grasp, the value of ε may be sensitive or insensitive to such perturbations. If the perturbations are "large," the origin will leave \mathcal{C}_w , so we cannot certify force closure. Let w_l and \bar{w}_l denote elements of \mathcal{C}_w and $\mathcal{C}_{\bar{w}}$ respectively. The following result bounds the size of tolerable perturbations in terms of $\mathcal{C}_{\bar{w}}$.

Theorem 1. *If $w_l - \bar{w}_l \in -\mathcal{C}_{\bar{w}}$ for all $l \in \mathcal{L}$, then $0 \in \mathcal{C}_w$.*

Proof. Suppose for the sake of contradiction that $0 \notin \mathcal{C}_w$. Then, there exists a such that $a^\top w > 0$ for all $w \in \mathcal{C}_w$ by the separating hyperplane theorem. Further, there must exist some $k \in \mathcal{L}$ satisfying $a^\top \bar{w}_k \leq a^\top \bar{w}_l$ for all $l \in \mathcal{L}$.

Since $w_l - \bar{w}_l \in -\mathcal{C}_{\bar{w}}$ for each l , there exists a set of convex weights $\alpha \in \mathbb{R}^{n_w}$ (with $\alpha \succeq 0$, $\sum_{l=1}^{n_w} \alpha_l = 1$) such that $w_k - \bar{w}_k = -\sum_{l=1}^{n_w} \alpha_l \bar{w}_l$. Thus,

$$w_k = \bar{w}_k + (w_k - \bar{w}_k) = \bar{w}_k - \sum_{l=1}^{n_w} \alpha_l \bar{w}_l, \quad (1)$$

which implies

$$\begin{aligned} a^\top w_k &= a^\top \bar{w}_k - \sum_{l=1}^{n_w} \alpha_l a^\top \bar{w}_l \\ &\leq a^\top \bar{w}_k - a^\top \bar{w}_k \sum_{l=1}^{n_w} \alpha_l = 0. \end{aligned} \quad (2)$$

But, $a^\top w_k > 0$ since $w_k \in \mathcal{C}_w$, contradicting (2). \square

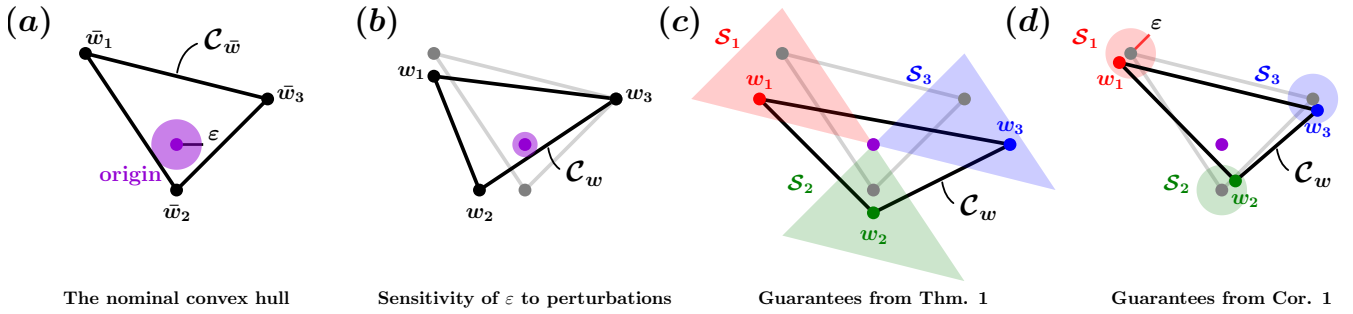


Fig. 2: (a) A set of nominal basis wrenches and their hull $\mathcal{C}_{\bar{w}}$ with a Ferrari-Canny metric of ε . (b) The value of ε is sensitive to some perturbations (e.g., w_2) but not others (e.g., w_1). We derive guarantees on the size of allowable perturbations such that $0 \in \mathcal{C}_w$. (c) The guarantees from Theorem 1. The shaded regions \mathcal{S}_i indicate areas where w_i may lie, independent of other basis wrenches, while still guaranteeing $0 \in \mathcal{C}_w$. Note how the \mathcal{S}_i are anisotropic. (d) The Ferrari-Canny metric ε also provides (strictly worse) guarantees as a result. Note this cartoon shows a planar wrench space, which is actually \mathbb{R}^6 .

We view $w_l - \bar{w}_l$ as the deviation of the l^{th} basis wrench from some nominal value. If all deviations are respectively contained in $-\mathcal{C}_{\bar{w}}$, then by Theorem 1, the origin lies in the convex hull \mathcal{C}_w of the true, perturbed basis wrenches, i.e., the grasp is force closure. See Fig. 2c for visual intuition.

While force closure is a binary measure of grasp quality, the Ferrari-Canny metric [5] reports the (extrinsic) robustness of force closure grasps by computing the radius ε of the largest origin-centered ball contained in \mathcal{C}_w . The grasp can resist any applied disturbance wrench with a norm smaller than ε , with the type of norm depending on assumed constraints on a hand’s control authority [5].

Corollary 1. *Let ε denote the Ferrari-Canny metric for a nominal set of basis wrenches $\bar{\mathcal{W}}$. If the true (i.e., perturbed) basis wrenches w_1, \dots, w_{n_w} satisfy*

$$w_l \in B_\varepsilon(\bar{w}_l), \forall l \in \mathcal{L}, \quad (3)$$

then the true grasp is force closure.

Proof. Follows by Theorem 1, since $B_\varepsilon(0) \subset -\mathcal{C}_{\bar{w}}$. Therefore, $0 \in \mathcal{C}_w$, so the true grasp is force closure. \square

Though ε represents a grasp’s sensitivity to external disturbance wrenches, Corollary 1 shows that ε also captures a grasp’s intrinsic robustness (e.g., to object pose/geometry) implicitly through friction cone uncertainty. We analyze concrete ways to explicitly use the relation between object geometry and basis wrenches in Sec. IV-A. See Fig. 2d for an illustration of Corollary 1.

B. Bounding ε with the Min-Weight Metric ℓ^*

One challenge for grasp synthesis with the Ferrari-Canny metric ε is that it is slow to optimize, especially due to its non-differentiability [6]. This motivates efficient approximations of ε like the min-weight metric ℓ^* , an almost-everywhere differentiable proxy for the Ferrari-Canny metric computable by solving a simple LP. In this section, we leverage Corollary 1 (and by extension, Theorem 1) to demonstrate that maximizing ℓ^* provides guarantees on intrinsic robustness. The key heuristic of the min-weight metric is to encourage the origin to lie far inside \mathcal{C}_w by maximizing its minimum convex weight (see [6]).

The following LP defines ℓ^* :

$$\begin{aligned}
 \ell^* = \text{maximize} \quad & \ell \\
 & \alpha \in \mathbb{R}^{n_w}, \ell \in \mathbb{R} \\
 \text{subject to} \quad & W\alpha = 0 \\
 & \mathbf{1}_{n_w}^\top \alpha = 1 \\
 & \alpha \succeq \ell \mathbf{1}_{n_w},
 \end{aligned} \quad (4)$$

where $W = W(q)$ is the matrix whose l^{th} column is $w_l \in \mathcal{W}$ and $\mathbf{1}_{n_w} \in \mathbb{R}^{n_w}$ is the vector of all 1s.

In [6, Fig. 3], it was observed that when empirically evaluated on force closure grasps, (an unknown positive scaling K of) ℓ^* seemed to lower bound ε uniformly. This motivated the maximization of ℓ^* as a conservative approximation of ε . If true, this would imply that ℓ^* captures intrinsic robustness by Corollary 1, since $B_{K\ell^*}(0) \subseteq B_\varepsilon(0)$. In this section, we formally prove the claim for “sufficiently regular” grasps.

Every compact set \mathcal{X} has a largest (not necessarily origin-centered) inscribing ball whose radius is called the Chebyshev radius, which we denote δ . We use this to impose some desired regularity on the types of grasps we consider.

Definition 1. *Let $P(\mathcal{X})$ denote the power set of \mathcal{X} . The set $\mathcal{S}_{\mathcal{O}} \subset P(\mathbb{R}^6)$ is the set of all sets of basis wrenches \mathcal{W} with cardinality n_w such that*

- 1) \mathcal{W} is possible to generate on the manipuland \mathcal{O} , and
- 2) the Chebyshev radii of \mathcal{C}_w , denoted $\delta(\mathcal{W})$, are uniformly bounded: $\delta(\mathcal{W}) \geq \delta_{\min} > 0$ for all $\mathcal{W} \in \mathcal{S}_{\mathcal{O}}$.

Condition (2) is mild; for every grasp, $\delta \geq \varepsilon$, so Definition 1 excludes grasps that are poor by virtue of a small Chebyshev radius up to a fixed level δ_{\min} . This condition is satisfied in practice, since we synthesize grasps by lower bounding their desired quality. This gives the following.

Theorem 2. *For all basis wrenches $\mathcal{W} \in \mathcal{S}_{\mathcal{O}}$ on a compact manipuland \mathcal{O} , there exists a grasp-independent constant $K(\mathcal{O}) > 0$ uniformly satisfying $K\ell^*(\mathcal{W}) \leq \varepsilon(\mathcal{W})$.*

C. Proof of Theorem 2

To prove Theorem 2, we require some intermediate results. Some details are elided to this paper’s extended version [21] for brevity. First, consider the following optimization program, whose optimal value is the Ferrari-Canny metric ε .



Fig. 3: **Left.** Hardware setup. We synthesized grasps on eight objects for an Allegro hand mounted on a Franka Research 3. From left to right: 3D-printed part, goblet, box, mug, cube, bottle, conditioner, and apple. We employed an eye-in-hand setup with a Zed camera to capture images used to train a NeRF prior to grasping. Only monocular RGB was used. **Right.** Representative FRoGGeR-synthesized grasps.

Lemma 1. When $0 \in \mathcal{C}_w$ (i.e., the grasp is force closure),

$$\begin{aligned} \varepsilon = \underset{a \in \mathbb{R}^6, b \in \mathbb{R}}{\text{minimize}} \quad & b \\ \text{subject to} \quad & a^\top w_l + b \geq 0, \forall l \in \mathcal{L}, \\ & a^\top a = 1. \end{aligned} \quad (5)$$

Proof. See the extended version [21, App. A]. \square

We compare the optimal value of (5) to ℓ^* by analyzing the dual of the min-weight LP. By a routine calculation of the dual of (4) (see [21, App. B]) and strong duality, we have

$$\begin{aligned} \ell^* = \underset{\nu \in \mathbb{R}^6, \varphi \in \mathbb{R}}{\text{minimize}} \quad & \varphi \\ \text{subject to} \quad & \nu^\top w_l + \varphi \geq 0, \forall l \in \mathcal{L} \\ & \sum_l (\nu^\top w_l + \varphi) = 1. \end{aligned} \quad (6)$$

We are now ready to prove Theorem 2.

Proof. We only consider the case where $0 \in \mathcal{C}_w$, since otherwise, there is no origin-centered inscribed ball of \mathcal{C}_w . This allows us to invoke Lemma 1. Consider an arbitrary set of basis wrenches $\mathcal{W} \in \mathcal{S}_O$ and suppose (a^*, b^*) is an optimal solution to (5) (so $\|a^*\|_2 = 1$). The pair (ν, φ) defined

$$\nu = \frac{a^*}{\sum_l (w_l^\top a^* + b^*)}, \quad \varphi = \frac{b^*}{\sum_l (w_l^\top a^* + b^*)} \quad (7)$$

is clearly feasible for (6). Dividing by $\sum_l (w_l^\top a^* + b^*) \geq 0$ is well-defined, since if the sum were not strictly positive, it would imply that \mathcal{C}_w is degenerate and contained in the hyperplane $\mathcal{P} = \{x \mid x^\top a^* + b^* = 0\}$, implying $\mathcal{W} \notin \mathcal{S}_O$.

By strong duality and the feasibility¹ of (ν, φ) , we have

$$\ell^* = \varphi^* \leq \frac{b^*}{\sum_l (w_l^\top a^* + b^*)} = \frac{\varepsilon}{\sum_l (w_l^\top a^* + b^*)}, \quad (8)$$

where $b^* = \varepsilon$ by Lemma 1. So, if $K > 0$ exists such that

$$K \leq \sum_l (w_l^\top a^*(\mathcal{W}) + b^*(\mathcal{W})) \quad (9)$$

uniformly for all $\mathcal{W} \in \mathcal{S}_O$, the proof follows.

Let $\delta(\mathcal{W}) > 0$ be the Chebyshev radius of \mathcal{C}_w . We show if $v \in \mathcal{C}_w$ is a point in \mathcal{C}_w farthest from \mathcal{P} , then $v^\top a^* + b^* \geq 2\delta$.

¹Since the primal (4) is a maximization, the dual (6) is a minimization and dual-feasible values *upper bound* the optimal dual and primal values.

Suppose that $v^\top a^* + b^* < 2\delta$. Let c be a Chebyshev center. Then, the distance of c to \mathcal{P} satisfies $c^\top a^* + b^* \geq \delta$, or c could not be the center of a ball inscribing \mathcal{C}_w . Let

$$y := c + \delta a^* \in B_\delta(c) \subset \mathcal{C}_w. \quad (10)$$

The distance of y to plane \mathcal{P} is given by

$$\begin{aligned} y^\top a^* + b^* &= c^\top a^* + \delta \cdot (a^*)^\top a^* + b^* \\ &= \delta + c^\top a^* + b^* \geq 2\delta > a^\top v + b, \end{aligned} \quad (11)$$

where the last inequality is by assumption, contradicting that v is a farthest point in \mathcal{C}_w from \mathcal{P} ; thus, $v^\top a^* + b^* \geq 2\delta$.

Now, there exists an extreme point of \mathcal{C}_w that is a farthest point in \mathcal{C}_w from \mathcal{P} , and \mathcal{W} contains all extreme points of \mathcal{C}_w by construction. Thus, there exists $w \in \mathcal{W}$ such that $w^\top a^* + b^* \geq 2\delta$. By Def. 1, we assume a δ_{\min} such that

$$\begin{aligned} 2\delta_{\min} \leq 2\delta(\mathcal{W}) &\leq \max_{w \in \mathcal{W}} w^\top a^*(\mathcal{W}) + b^*(\mathcal{W}) \\ &\leq \sum_l (w_l^\top a^*(\mathcal{W}) + b^*(\mathcal{W})), \end{aligned} \quad (12)$$

so the choice $K = 2\delta_{\min}$ completes the proof. \square

III. HARDWARE EXPERIMENTS

We now compare real-world grasps synthesized by maximizing the min-weight metric vs. the method of Wu *et al.* [18], which simply finds feasible force closure grasps. Prior work has shown that min-weight synthesized grasps outperform Wu's method under simulated shaking [6], showing that ℓ^* is robust to extrinsic uncertainties. Here, we evaluate *intrinsic* robustness to unknown uncertainties arising from, e.g., poor object reconstruction. This provides some preliminary evidence of our theory's usefulness. We leave analysis and testing of other grasp metrics for future work.

A. Grasp Synthesis with Bilevel Nonlinear Optimization

Both methods synthesize grasps via nonlinear gradient-based optimization. In [6], the min-weight-based program is called *FRoGGeR*, and has the following form:

$$\begin{aligned} &\underset{q}{\text{maximize}} \quad \ell^*(q) \\ &\text{subject to} \quad q_{\min} \leq q \leq q_{\max} \\ & \quad \ell^*(q) \geq k_\ell/n_w \\ & \quad FK^i(q) \in \partial\mathcal{O}, \quad i = 1, \dots, n_f \\ & \quad \text{No (non-finger/object) collision.} \end{aligned} \quad (\text{FRoGGeR})$$

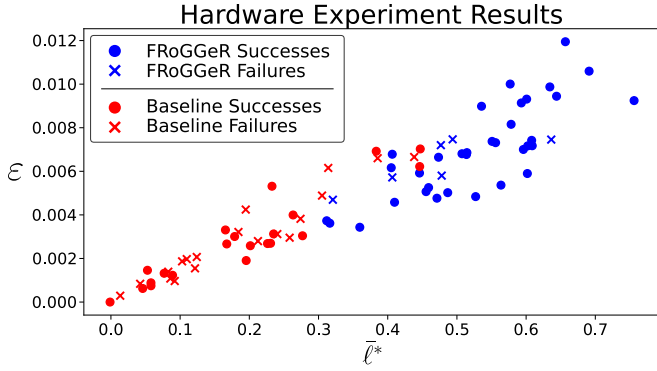


Fig. 4: Hardware results. FROGGER-synthesized grasps are of higher quality and have a much higher success rate (34/40) than the baseline, which only ensures force closure and does not optimize for grasp quality, resulting in many low-quality grasps and a lower success rate (21/40). This plot also confirms the validity of Theorem 2, as a linear lower bound can be seen.

Wu’s method [18] instead solves a feasibility program:

$$\begin{aligned}
 & \text{find } q \\
 & \text{subject to } q_{\min} \preceq q \preceq q_{\max} \\
 & J(q) = 0 \\
 & FK^i(q) \in \partial\mathcal{O}, i = 1, \dots, n_f \\
 & \text{No (non-finger/object) collision,}
 \end{aligned} \tag{Wu}$$

where $J(q)$ is the value of an optimization program that equals 0 when a grasp is force closure [18, Eqn. 5].

Both (FROGGER) and (Wu) are bilevel optimization programs, since $\ell^*(q)$ and $J(q)$ are themselves optimal values of lower-level programs. However, ℓ^* is a measure of a grasp’s robustness, while J only indicates whether a grasp is force closure. Thus, (FROGGER) also maximizes a grasp’s quality, which we show is vital to real-world performance. The last two constraints are handled using implicit functions by letting the surface $\partial\mathcal{O}$ be the 0-level set of a smooth function $s : \mathbb{R}^3 \rightarrow \mathbb{R}$. For details, see [18], [6].

B. Experimental Results

For both methods, we synthesized 5 grasps on 8 objects (see Fig. 3) for a total of 40 grasps each, comparing the success/failure rates as well as the values of *normalized min-weight metric* $\bar{\ell}^* \in [0, 1]$ (where $\bar{\ell}^* := \ell^*/n_w$) and ϵ . To synthesize grasps using these methods, some object representation is required. We chose to train a NeRF [22] of the object by scanning the scene with a wrist-mounted camera. Then, a mesh was extracted from the NeRF by running the marching cubes algorithm on a density level set, which we chose empirically to be 17.5 for all objects.

Unlike Wu *et al.*, we did not use a learned model to extract initial guesses for the nonlinear optimization program. Instead, we used the heuristic sampler described in [6] for both methods, and the optimization program we solved was identical except that (a) (Wu) is a feasibility program, so there is no objective, and (b) (Wu) enforces a force closure equality constraint, while (FROGGER) enforces that $\bar{\ell}^* \geq k_\ell$, where we used $k_\ell = 0.3$ as in [6].

We used the following simple controller: for an optimal solution q^* of (FROGGER) or (Wu), we have corresponding

fingertip positions $x^i = FK^i(q^*)$. We defined new target positions for each finger 2cm into the surface,

$$x_{\text{new}}^i = x^i + 0.02\nabla s(x^i), \tag{13}$$

and let q_{new}^* satisfy $x_{\text{new}}^i = FK^i(q_{\text{new}}^*), \forall i \in \mathcal{I}$, computed using inverse kinematics and tracked with P control:

$$\tau_{\text{fb}} = -K_p(q - q_{\text{new}}^*). \tag{14}$$

After grasping, the object was picked up 20cm and held in place. A grasp was said to fail if it was not successfully picked or if after picking, there was significant rotation.

The quantitative results are summarized in Fig. 4. Overall, grasps synthesized with (FROGGER) succeeded in 34/40 trials (85%), while grasps synthesized using the baseline (Wu) succeeded in 21/40 (52.5%). Baseline grasps tended to have low values of ϵ and $\bar{\ell}^*$, which supports the theory of uncertainty awareness in grasps detailed in Sec. II. Further, we see that the relation between $\bar{\ell}^*$ and ϵ supports Theorem 2, as a linear lower bound is visually clear.

C. Discussion

Due to perception error/uncertainty in the NeRF training, the pose and geometry of the object was sometimes inaccurate. This led to some trials where during the pre-grasp motion, the object was struck by the manipulator prior to grasping. These trials were excluded ($\sim 15\%$ of attempts), since the purpose of the experiments was to evaluate the robustness of grasps themselves. Qualitatively, the dominant failure mode for both methods was attempting grasps on regions where grasp quality was sensitive to shape error. For instance, two failures for FROGGER grasps were caused by trying to grab the tea bottle on the sloped area near the top, and many of the failures for Wu’s method were caused by trying to grab edges of objects and slipping.

Finally, we note that our reported success rates for Wu’s method are significantly lower than those reported in [18]. We suspect that one cause for this is that the initial guesses supplied by Wu’s learning-based model are critical for generating good grasps. Additionally, our grasping controller is only position-based, whereas in [18], they command the forces found by computing $J(q)$. The discrepancy could also be due to the different hardware setup, slight differences in implementation, or randomness. Despite these differences, FROGGER-synthesized grasps are able to achieve about the same success rate without a learning-based model.

IV. TOWARDS PROBABILISTIC FORCE CLOSURE

The previous analysis applied the framework of intrinsic robustness and Theorem 1 to characterize the min-weight metric. We now explore another application, and develop the idea of *probabilistic force closure*, wherein we use the relation between object geometry and basis wrenches to apply our theory explicitly when we have a *belief distribution* over the object geometry. We assume we have such a distribution, though obtaining one in practice is challenging and the general case is left for future work.

Suppose the basis wrenches \mathcal{W} associated with a grasp are a random variable generated by, e.g., uncertainty in the contact locations x^i , the surface normals $n^i(x^i)$ of \mathcal{O} , parameters γ of the object $\mathcal{O}(\gamma)$, etc. Then, one can attempt to compute or lower bound the probability of force closure.

Let $\bar{\mathcal{W}}$ and \mathcal{W} be nominal and random sets of basis wrenches respectively with convex hulls $\mathcal{C}_{\bar{w}}$ and \mathcal{C}_w . By Theorem 1, we can bound the probability of force closure:

$$\begin{aligned} \mathbb{P}[w_l - \bar{w}_l \in -\mathcal{C}_{\bar{w}}, \forall l \in \mathcal{L}] &\leq \mathbb{P}[0 \in \mathcal{C}_w] \\ &\leq \mathbb{P}[\text{force closure}]. \end{aligned} \quad (15)$$

If we can compute or bound the LHS of (15), then probabilistically-robust grasps may be synthesized. However, as-is, (15) is not very useful, since (a) it requires a probability measure to be defined over \mathcal{O} , and (b) the relevant uncertainty sets (e.g., $-\mathcal{C}_{\bar{w}}$ in Theorem 1 or $B_\varepsilon(\bar{w}_l)$ in Corollary 1) must be easily integrable with respect to this measure.

A. Probabilistic Object Normals for Grasping

In light of these challenges, we study a special case where the contact locations x^i are known but the surface normals n^i are random and modeled by a degenerate Gaussian distribution. This gives rise to a method we call *probabilistic object normals for grasping* (PONG).

Let the known contacts be denoted $\mathcal{X} = \{x^1, \dots, x^{n_f}\}$ and consider the set of normals that induce basis wrenches certifying force closure. We call this the *force closure set*:

$$\mathfrak{N}(\mathcal{X}) := \{\{n^i\}_{i=1}^{n_f} \mid 0 \in \mathcal{C}_w\}. \quad (16)$$

Suppose the normals are jointly distributed with some density function $p(n^1, \dots, n^{n_f})$ and that they are mutually independent such that p can be factorized as $\prod_{i=1}^{n_f} p(n^i)$. Then, we can compute the probability that $0 \in \mathcal{C}_w$:

$$\mathbb{P}[0 \in \mathcal{C}_w] = \int_{\mathfrak{N}(\mathcal{X})} \prod_{i=1}^{n_f} p(n^i) dn^i. \quad (17)$$

Though p is factorizable, the integral itself is not, since the integration variables are coupled by the domain \mathfrak{N} , and \mathcal{C}_w depends on all of the random normals. Further, \mathfrak{N} is difficult to parameterize; it has no closed form since it is implicitly defined by the condition $0 \in \mathcal{C}_w$, where hull membership is checked by solving an LP [6].

Thus, our strategy is to derive an *approximate force closure set* $\mathcal{A} \subseteq \mathfrak{N}$ in tandem with a choice of density function p for which the integral over \mathcal{A} is known. Since $\mathcal{A} \subseteq \mathfrak{N}$, integrating p over \mathcal{A} yields a lower bound on (17).

B. Random Normals, Forces, and Wrenches

We model each random normal n^i as the sum of a deterministic *mean normal vector* denoted \bar{n}^i and a *random perturbation vector* $\Delta n^i = \mathcal{T}_i z$, where $\mathcal{T}_i = \begin{bmatrix} \bar{t}_1^i & \bar{t}_2^i \end{bmatrix} \in \mathbb{R}^{3 \times 2}$ is a basis for the tangent plane at \bar{n}^i , and $z \sim \mathcal{N}(0, \Sigma^i)$ is a zero-mean Gaussian random vector in \mathbb{R}^2 .

We parameterize the normals in this way for two reasons. First, all perturbations in the direction of \bar{n}^i do not change its direction, and thus do not contribute to uncertainty in

its orientation. Second, the planar restriction allows us to integrate a Gaussian density in the plane over a polytope using Green's Theorem [23, Prop. 1].

We remark that with this construction, the random normals n^i will not have unit norm. However, since friction cones are invariant under scaling, as long as the random basis forces f_j^i remain Coulomb-compliant, we can still certify force closure with the induced random basis wrenches.

We now construct a random friction pyramid about the random normal n^i such that its edges are always Coulomb-compliant. First, consider the *mean* basis wrenches \bar{w}_j^i with force and torque components \bar{f}_j^i and $\bar{\tau}_j^i$. We represent \bar{f}_j^i as the sum of \bar{n}^i and a tangent component $(\bar{f}_j^i)_t$ of length μ . We compute $(\bar{f}_j^i)_t$ using a unit length *generator* $\bar{g}_j^i(\bar{n}^i) \in \mathbb{R}^3$ of our choice orthogonal to \bar{n}^i such that

$$(\bar{f}_j^i)_t = \mu \cdot (\bar{g}_j^i \times \bar{n}^i). \quad (18)$$

For example, for some arbitrary $v \neq \bar{n}^i$, we could pick $\bar{g}_j^i(\bar{n}^i) = (v \times \bar{n}^i) / \|v \times \bar{n}^i\|_2$. Thus, we can write

$$\begin{aligned} \bar{f}_j^i &= \bar{n}^i + \mu \cdot (\bar{g}_j^i \times \bar{n}^i) \\ \implies \bar{w}_j^i &= \underbrace{\begin{bmatrix} (I + \mu \cdot [\bar{g}_j^i]_\times) \\ [x^i]_\times (I + \mu \cdot [\bar{g}_j^i]_\times) \end{bmatrix}}_{:=T_j^i(\bar{n}^i)} \bar{n}^i, \end{aligned} \quad (19)$$

and similarly, $w_j^i = T_j^i n^i$, which satisfies $\|(f_j^i)_t\|_2 \leq \mu \|n^i\|$ (in contrast, recall that $\|(\bar{f}_j^i)_t\|_2 = \mu \|n^i\|$). To see this, let ψ_j^i be the angle between \bar{g}_j^i and n^i . Then,

$$\|(f_j^i)_t\|_2 = \mu \|\bar{g}_j^i\| \|n^i\| \sin(\psi_j^i) \leq \mu \|n^i\|, \quad (20)$$

ensuring Coulomb-compliance of the random basis forces.

C. Deriving an Approximate Force Closure Set

Next, we define disjoint sets \mathcal{A}^i whose union defines a set $\mathcal{A} \subseteq \mathfrak{N}$. Combining Theorem 1 with the definition of \mathfrak{N} and the relation $w_j^i = T_j^i n^i$ from Sec. IV-B, if \mathcal{A}^i satisfies

$$\mathcal{A}^i \subseteq \{n^i \mid T_j^i(n^i - \bar{n}^i) \in -\mathcal{C}_{\bar{w}}, \forall j \in \mathcal{J}\}, \quad (21)$$

then we have

$$n^i \in \mathcal{A}^i, \forall i \in \mathcal{I} \implies (n^1, \dots, n^{n_f}) \in \mathfrak{N}. \quad (22)$$

Letting $\mathcal{A} = \bigcup_{i=1}^{n_f} \mathcal{A}^i$, we therefore have the bound

$$\begin{aligned} \prod_{i=1}^{n_f} \int_{\mathcal{A}^i} p(n^i) dn^i &= \int_{\mathcal{A}} \prod_{i=1}^{n_f} p(n^i) dn^i \\ &\leq \int_{\mathfrak{N}} \prod_{i=1}^{n_f} p(n^i) dn^i. \end{aligned} \quad (23)$$

We can factorize the integral this way since $\mathcal{C}_{\bar{w}}$ in (21) depends only on the *mean* normals $(\bar{n}^1, \dots, \bar{n}^{n_f})$, which are fixed parameters of the known uncertainty distributions. In contrast, in (16), \mathcal{C}_w (and thus \mathfrak{N}) depends jointly on all of the *random* normals, which are the integration variables.




Fig. 5: More examples of grasps on a synthetically-uncertain spherical manipulant. The surface normal tangent covariance at each point on the sphere is (proportional to) an isotropic monotonically-increasing function of the real part of the (2,4)-spherical harmonic. The synthesized grasps tend to avoid placing the fingertips on the red uncertain regions while remaining kinematically feasible. They were synthesized in 6.47 and 4.80 seconds respectively.

Unfortunately, the set on the RHS of (21) still cannot be expressed in closed form. Thus, we let \mathcal{A}^i conservatively approximate it by computing a polytope with large volume satisfying (21) that is parameterized by its vertices.

To do this, for each finger, we fix k search directions $d^{i,k} \in \mathbb{R}^3$ in the tangent plane to \mathcal{O} at x^i with index set \mathcal{K} . We compute the longest step $\theta^{i,k} \geq 0$ that can be taken in this direction while satisfying (21) and denote the point $\theta^{i,k}d^{i,k}$ as $v^{i,k}$. We then let \mathcal{A}^i be the convex hull over all $v^{i,k}$.

Conveniently, this search can be expressed as a linear program. Define $\bar{W} \in \mathbb{R}^{6 \times n_w}$, whose columns are the nominal basis wrenches. Then, to compute each $v^{i,k}$, we can solve the following LP for each $(i, k) \in \mathcal{I} \times \mathcal{K}$ in parallel:

$$\begin{aligned} & \text{maximize} && \theta^{i,k} && (24a) \\ & \theta^{i,k} \in \mathbb{R} && && \\ & \{\alpha_j^{i,k}\}_{j=1}^{n_s} \subset \mathbb{R}^{n_w} && && \end{aligned}$$

$$\text{subject to} \quad \theta^{i,k} \geq 0 \quad (24b)$$

$$\alpha_j^{i,k} \geq 0, \forall j \in \mathcal{J} \quad (24c)$$

$$\mathbb{1}_{n_w}^\top \alpha_j^{i,k} = 1, \forall j \in \mathcal{J} \quad (24d)$$

$$T_j^i(\theta^{i,k}d^{i,k}) = -\bar{W}\alpha_j^{i,k}, \forall j \in \mathcal{J}. \quad (24e)$$

Constraint (24b) enforces nonnegativity of the scaling along $d^{i,k}$, constraints (24c) and (24d) enforce that the $\alpha_j^{i,k}$ are valid convex weights, and constraint (24e) enforces that the random normal $n^i = \bar{n}^i + \theta^{i,k}d^{i,k}$ must lie in \mathcal{A}^i .

Therefore, after solving (24), we can compute

$$L_{FC} := \prod_{i=1}^{n_f} \int_{\mathcal{A}^i(\theta, d)} p(n^i) dn^i \quad (25)$$

using [23, Prop. 1], which bounds the probability of force closure. Due to space constraints, computational considerations for computing (25) are detailed in [21, App. C-E].

D. Toy Examples

We first demonstrate PONG on two toy examples. In both, we let the manipulant \mathcal{O} be a sphere of radius 0.05m, and we define (unnormalized) distributions over $x \in \partial\mathcal{O}$ as follows:

$$\begin{aligned} \bar{n}(x) &= x / \|x\|_2, \\ \sigma_1^2 &= 100x_3^2, \quad \sigma_2^2 = 0.01 \exp(Y_4^2(\theta(x), \phi(x))), \end{aligned} \quad (26)$$

where the mean of both is $\bar{n}(x)$ and they have isotropic tangent variances defined by σ_1^2 and σ_2^2 respectively.

The first distribution assigns low uncertainty at the equator of the object, symmetrically increasing it towards the poles. The second assigns uncertainties based on the (2,4)-spherical harmonic Y_4^2 , where $\theta(x), \phi(x)$ are the azimuthal and polar angles of $x \in \partial\mathcal{O}$. Visualizations of grasps generated with each are shown in Fig. 1 and Fig. 5 respectively.

We note the qualitative difference that the uncertainty distribution makes, even though the manipulant's geometry is identical in each case. This demonstrates that accounting for uncertainty in grasp synthesis can yield vastly different solutions: for example, a vision system could report higher uncertainty for regions of the object that are occluded, which in turn could encourage grasping visible regions.

E. Comparative Analysis on Realistic Meshes

While PONG-synthesized grasps on toy examples showed some promise, to use it on more realistic examples, we need a distribution supported on $\partial\mathcal{O}$. In the absence of a natural choice, we define a *curvature-regularized* one. At $x \in \partial\mathcal{O}$, the *shape operator* $\mathfrak{S}_x : T_x\partial\mathcal{O} \rightarrow T_x\partial\mathcal{O}$ is defined as the directional derivative $\mathfrak{S}_x(v) := -D_v N(x)$, where N is the unit normal at x . The eigenpairs of \mathfrak{S}_x , $(\kappa_1, v_1), (\kappa_2, v_2) \in \mathbb{R} \times T_x\partial\mathcal{O}$, are the *principal curvatures and directions* at x respectively [24].

Since $N(x) = \nabla s(x) / \|\nabla s(x)\|_2$, we have

$$\mathfrak{S}_x(v) = - \left(I - \frac{\nabla s(x)\nabla s(x)^\top}{\|\nabla s(x)\|_2^2} \right) \frac{\nabla^2 s(x)}{\|\nabla s(x)\|_2} v. \quad (27)$$

Letting the directions v_1, v_2 form our tangent basis at each contact x^i and (some increasing function of) the magnitude of κ_1, κ_2 form our variances, we can define a curvature-sensitive uncertainty distribution with parameters

$$\begin{aligned} \bar{n}^i &:= - [\nabla_x s(x^i) / \|\nabla_x s(x^i)\|_2], \\ \bar{t}_m^i &:= v_m, \quad m \in \{1, 2\}, \\ (\sigma_m^i)^2 &:= \log(K_{\text{curv}} \cdot |\kappa_m| + h), \quad m \in \{1, 2\}, \end{aligned} \quad (28)$$

where $K_{\text{curv}} > 0$ relates curvature to uncertainty and $h > 0$ captures prior uncertainty at x . Intuitively, we assign greater uncertainty to high curvatures, since slight errors in position will lead to large changes in the surface normal.

Since ℓ^* captures intrinsic uncertainty, we study whether it performs comparably to PONG even though it only provides

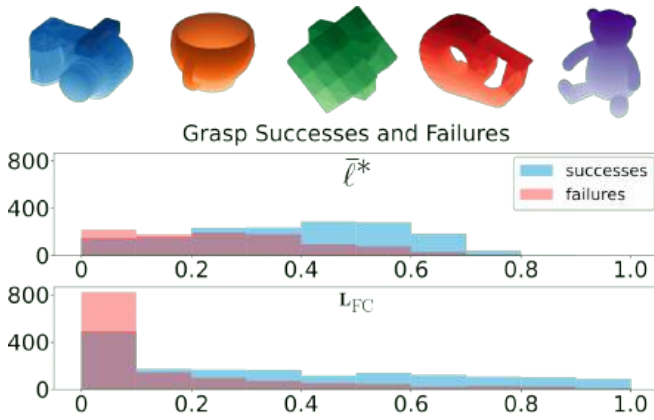


Fig. 6: **Top.** 5 simulated objects. **Bottom.** Success/failure histograms for 2000 simulated grasps for ℓ^* and L_{FC} .

conservative guarantees on uncertainty-awareness per Corollary 1 without reasoning about the bound (15), and find that it does. We synthesize 2000 grasps on 5 objects (a camera, teacup, rotated Rubik’s cube, tape dispenser, and teddy bear) and in simulation study pick success as a function of metric.

All 2000 grasps are synthesized using (FRoGGeR) but with objective $L_{FC}(q)$ and evaluated on ℓ^* and L_{FC} . We observe two key trends (see Fig. 6). Based on the theory of Sec. II, we expect fewer failures as each metric rises, which we observe. Further, because each metric is conservative, we expect the ratio of successes to failures to rise with the metric, which is again confirmed. This provides evidence that the min-weight metric is empirically probabilistically sensitive to certain uncertainties in object geometry.

V. CONCLUSION

This work develops a theory of intrinsic uncertainty and robustness for dexterous grasps. Theorems 1 and 2 provide guarantees on the intrinsic robustness of ℓ^* -synthesized grasps, motivating its usage as a principled and efficient metric. These guarantees were tested on hardware against a competitive baseline, demonstrating the importance of robustness in grasp synthesis. We also use our theory to develop PONG, a method for probabilistically-aware grasp synthesis. We show that PONG can generate grasps that account for uncertainty over the object surface, paving the way for using stochastic models of geometric uncertainty with grasp planners. In future work, we will study acquiring uncertainty distributions compatible with PONG to more explicitly account for intrinsic uncertainty in dexterous grasps.

REFERENCES

- [1] Máximo A. Roa and Raúl Suárez. Grasp quality measures: review and performance. *Autonomous Robots*, 38:65 – 88, 2014.
- [2] Z. Li and S.S. Sastry. Task-oriented optimal grasping by multifingered robot hands. *IEEE Journal on Robotics and Automation*, 4(1):32–44, 1988.
- [3] Daniel Kappler, Jeannette Bohg, and Stefan Schaal. Leveraging big data for grasp planning. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4304–4311, 2015.
- [4] Lin Shao, Fábio Ferreira, Mikael Jorda, Varun Nambiar, Jianlan Luo, Eugen Solowjow, Juan Aparicio Ojea, Oussama Khatib, and Jeannette Bohg. UniGrasp: Learning a Unified Model to Grasp with N-Fingered Robotic Hands. *ArXiv*, abs/1910.10900, 2019.

- [5] C. Ferrari and J. Canny. Planning optimal grasps. In *Proceedings 1992 IEEE International Conference on Robotics and Automation*, pages 2290–2295 vol.3, 1992.
- [6] Albert H. Li, Preston Culbertson, Joel W. Burdick, and Aaron D. Ames. FRoGGeR: Fast Robust Grasp Generation via the Min-Weight Metric. *ArXiv*, abs/2302.13687, 2023.
- [7] Jonathan Weisz and Peter K. Allen. Pose error robust grasping from contact wrench space metrics. In *2012 IEEE International Conference on Robotics and Automation*, pages 557–562, 2012.
- [8] Stanimir Dragiev, Marc Toussaint, and Michael Gienger. Gaussian process implicit surfaces for shape estimation and grasping. In *2011 IEEE International Conference on Robotics and Automation*, pages 2845–2850, 2011.
- [9] Jeffrey Mahler, Sachin Patil, Ben Kehoe, Jur van den Berg, Matei Ciocarlie, Pieter Abbeel, and Ken Goldberg. GP-GPIS-OPT: Grasp planning with shape uncertainty using Gaussian process implicit surfaces and Sequential Convex Programming. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4919–4926, 2015.
- [10] Cristiana de Farias, Naresh Marturi, Rustam Stolkin, and Yasemin Bekiroglu. Simultaneous Tactile Exploration and Grasp Refinement for Unknown Objects. *IEEE Robotics and Automation Letters*, PP, 02 2021.
- [11] Miao Li, Kaiyu Hang, Danica Kragic, and Aude Billard. Dexterous grasping under shape uncertainty. *Robotics and Autonomous Systems*, 75:352–364, 2016.
- [12] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, and Ken Goldberg. Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics. *ArXiv*, abs/1703.09312, 2017.
- [13] Qingkai Lu, Kautilya Chenna, Balakumar Sundaralingam, and Tucker Hermans. Planning Multi-Fingered Grasps as Probabilistic Inference in a Learned Deep Network. In *International Symposium of Robotics Research*, 2018.
- [14] Qingkai Lu, Mark Van der Merwe, Balakumar Sundaralingam, and Tucker Hermans. Multifingered Grasp Planning via Inference in Deep Neural Networks: Outperforming Sampling by Learning Differentiable Models. *IEEE Robotics and Automation Magazine*, 27(2):55–65, 2020.
- [15] Min Liu, Zherong Pan, Kai Xu, Kanishka Ganguly, and Dinesh Manocha. Deep Differentiable Grasp Planner for High-DOF Grippers. *ArXiv*, abs/2002.01530, 2020.
- [16] Tengyu Liu, Zeyu Liu, Ziyuan Jiao, Yixin Zhu, and Song-Chun Zhu. Synthesizing Diverse and Physically Stable Grasps With Arbitrary Hand Structures Using Differentiable Force Closure Estimator. *IEEE Robotics and Automation Letters*, 7:470–477, 2021.
- [17] Ruicheng Wang, Jialiang Zhang, Jiayi Chen, Yinchen Xu, Puhao Li, Tengyu Liu, and He Wang. DexGraspNet: A Large-Scale Robotic Dexterous Grasp Dataset for General Objects Based on Simulation. *ArXiv*, abs/2210.02697, 2022.
- [18] Albert Wu, Michelle Guo, and C. Karen Liu. Learning Diverse and Physically Feasible Dexterous Grasps with Generative Model and Bilevel Optimization. *ArXiv*, abs/2207.00195, 2022.
- [19] Florian T. Pokorny and Danica Kragic. Classical grasp quality evaluation: New algorithms and theory. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3493–3500, 2013.
- [20] Hongkai Dai, Anirudha Majumdar, and Russ Tedrake. Synthesis and Optimization of Force Closure Grasps via Sequential Semidefinite Programming. In *International Symposium of Robotics Research*, 2015.
- [21] Albert H. Li, Preston Culbertson, and Aaron D. Ames. Toward An Analytic Theory of Intrinsic Robustness for Dexterous Grasping. *ArXiv*, abs/2403.07249, 2024.
- [22] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *CoRR*, abs/2003.08934, 2020.
- [23] Naoki Hayashi, Kohei Segawa, and Shigemasa Takai. 2D Voronoi Coverage Control with Gaussian Density Functions by Line Integration. *SICE Journal of Control, Measurement, and System Integration*, 10(2):110–116, 2017.
- [24] B. O’Neill. *Elementary Differential Geometry, Revised 2nd Edition*. Elsevier Science, 2006.