

IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping with the Incremental Potential Contact Model

Chung Min Kim



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2022-72

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2022/EECS-2022-72.html>

May 12, 2022

Copyright © 2022, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

Thank you to the following people who have contributed to the content of this thesis: Mike Danielczuk, Isabella Huang, and Professor Ken Goldberg. I additionally thank Jeffrey Ichnowski who introduced Incremental Potential Contact (IPC) to us, and Minchen Li and Zachary Ferguson for support in using IPC.

IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping with the
Incremental Potential Contact Model

by

Chung Min Kim

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Master of Science

in

Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ken Goldberg, Chair
Professor Sergey Levine

Spring 2022

**IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping
with the Incremental Potential Contact Model**

by Chung Min Kim

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, in partial satisfaction of the requirements for the
degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:



Professor Ken Goldberg
Research Advisor

11 May 2022

(Date)

* * * * *



Professor Sergey Levine
Second Reader

5/11/22

(Date)

IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping with the
Incremental Potential Contact Model

Copyright 2022
by
Chung Min Kim

Abstract

IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping with the
Incremental Potential Contact Model

by

Chung Min Kim

Master of Science in Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Ken Goldberg, Chair

Accurately simulating whether an object will be lifted securely or dropped during grasping is a longstanding Sim2Real challenge. Soft compliant jaw tips are almost universally used with parallel-jaw robot grippers due to their ability to increase contact area and friction between the jaws and the object to be manipulated. However, interactions between the compliant surfaces and rigid objects are notoriously difficult to model.

We introduce IPC-GraspSim, a novel grasp simulator that extends Incremental Potential Contact (IPC) — a highly accurate collision + deformation model developed in 2020 for computer graphics. IPC-GraspSim models both the dynamics and the deformation of compliant jaw tips to reduce Sim2Real gap for robot grasping. We evaluate IPC-GraspSim using a set of 2,000 physical grasps across 16 adversarial objects where analytic models perform poorly. In comparison to both analytic quasistatic contact models (soft point contact, REACH, 6DFC) and dynamic grasp simulators (Isaac Gym with FleX), results suggest IPC-GraspSim can predict robustness with higher precision and recall ($F1 = 0.85$). IPC-GraspSim increases F1 score by 0.03 to 0.20 over analytic baselines and 0.09 over Isaac Gym, at a cost of 8000x and 1.5x more compute time, respectively.

To my family,

Contents

Contents	ii
List of Figures	iii
List of Tables	v
1 Introduction	1
2 Related Work	4
2.1 Grasp Metrics and Contact Models for Soft Grippers	4
2.2 Grasp Simulators for Soft Grippers	5
2.3 Simulation of Deformable Materials	5
2.4 Incremental Potential Contact	6
3 IPC-GraspSim	8
3.1 Problem Statement	8
3.2 IPC-GraspSim	9
3.3 Parameter Sensitivity Analysis	12
4 Experiments	15
4.1 Experiment Setup	15
4.2 Benchmark Grasp Models	15
4.3 Failure Mode Analysis	17
5 Discussion	19
5.1 Conclusion	19
5.2 Future Work	19
5.3 Reflections	20
Bibliography	21

List of Figures

1.1	Simulated (left) and physical (right) grasp outcomes: in the first two rows, IPC-GraspSim correctly predicts grasp success and grasp failure. In the last two rows, predictions are incorrect due to unmodeled cantilever bending + rotations due to pushing.	3
2.1	Cube grasps with incorrect contact handling: (A) interpenetration, (B) force-field effect.	6
2.2	“Squishy ball test” figure from IPC [27]: a soft toy with many tendrils is thrown into a glass wall with high velocity. The custom collision detection implementation and contact modeling allows the microstructures to disentangle themselves post-impact with no interpenetration artifacts.	7
3.1	Two grasps on a cube with IPC-GraspSim, showing state at successive timesteps. (Top row) Successful grasp: 1) cube and gripper initialized; 2) cube contacts right jaw; 3) cube contacts left jaw; 4) gripper lifts cube. (Bottom row) Failed grasp: 1) cube and gripper initialized; 2) cube contacts right jaw; 3) cube <i>misses</i> left jaw; 4) jaws meet each other.	9
3.2	Corner grasp on cube: (left) gripper perturbed away from object, physical grasp attempt fails (right) gripper perturbed towards the object, physical grasp attempt succeeds. Both grasp trials intend to perform the same action, yet uncertainties allow the grasp success to vary across trials.	10
3.3	Tested values for Ψ_{th} , and the gripper pad deformation. Sides of the gripper pad are colored in red. $\Psi_{th} = 5 \times 10^4$ models the gripper pad thinning under pressure, without deforming in an unrealistic way.	12
3.4	The 21 adversarial objects used in the grasp dataset collected by Danielczuk et al. [8], where analytic grasp models perform poorly. Each object is sampled with 25 test grasps, 5 trials per grasp. The training object set (cat, pear, vase) for our experiments is highlighted in purple.	13

- 4.1 Four grasps (columns) with three simulators: (Top to bottom) Isaac Gym, IPC-GraspSim, Physical Experiments. (A) False negative: IPC-GraspSim and Isaac Gym fail to predict that the right jaw will contact the cube. (B) True positive: bowl correctly moves and rotates into gripper for both sims. (C) False positive: In Isaac Gym, incorrect contact forces push “pawn” up and away from the plane. In IPC-GraspSim, rigid jaws have sufficient contact area to grasp the object. Physical grasp fails as jaw cantilevering decreases contact area. (D) False negative: For Isaac Gym, the object is pushed upwards, with a similar issue as described in (C). In IPC-GraspSim, left gripper jaw pushes object out of the gripper. Videos of all simulator grasp attempts can be found on the website. 18

List of Tables

3.1	Rectangular jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 0.86.	14
3.2	Rounded jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 0.91.	14
4.1	Grasp robustness values predicted by IPC-GraspSim (top row) and observed from 125 physical experiments (bottom row) using an ABB YuMi with a compliant parallel-jaw gripper on the “vase” object shown in Fig. 1. Each column corresponds to a unique grasp configuration (i.e., relative pose between gripper and object). Each cell contains the robustness $R = \frac{1}{N} \sum_{i \in \{1..N\}} S_i$, with $N = 5$, and is colored to indicate its relative robustness (red: unreliable, green: robust). We observe a very high correlation of predicted grasp robustness between the two rows. Example true positive (grasp 2), true negative (grasp 0), false positive (grasp 7), and false negative (grasp 1) grasps are shown in Fig. 1. See Section 4.3 for discussion of grasp configurations 1 and 7.	16
4.2	Average Precision (AP) and Average Recall (AR) for each model’s grasp quality predictions for the 2000 grasps collected on the physical robot. All runtimes are provided per-grasp. Per-frame runtimes for Isaac Gym and IPC-GraspSim are 0.28 and 6.05 seconds respectively.	16

Acknowledgments

Thank you to the following people who have contributed to the content of this thesis: Mike Danielczuk, Isabella Huang, and Professor Ken Goldberg. I additionally thank Jeffrey Ichnowski who introduced Incremental Potential Contact (IPC) to us, and Minchen Li and Zachary Ferguson for support in using IPC.

Chapter 1

Introduction

Recent work in grasping objects with parallel-jaw grippers uses deep neural networks to estimate high-quality grasps directly from RGB or depth images [11, 18, 25, 33, 36, 38, 47]. These grasping networks are trained on large-scale datasets of labeled grasps generated either on physical systems [11, 18, 25, 38] or in simulation [17, 31, 33]. Simulation-based approaches can avoid time-intensive physical grasp attempts and have shown strong transfer to physical systems [33], but may suffer from noisy labels due to errors in modeling the interaction between gripper jaws and objects.

Sim2Real for grasping is complicated by soft materials added to the jaw tips on almost all parallel-jaw robot grippers, from soft foam padding [44] to custom-made silicone rubber fingertips [33]. Many simulators use discrete collision detection and penalty-based contact methods to determine contact forces at each timestep. Although computationally efficient, these simulators require that collision distances and timestep lengths be manually tuned for different object geometries and material parameters to avoid interpenetration or other physically unrealistic interactions between the jaws and object. Even so, object meshes may intersect during large timesteps or large contact force magnitudes, leading to false positive and false negative grasp predictions.

In 2020, Li et al. [27] introduced Incremental Potential Contact (IPC), a simulator that uses a novel collision model and continuous collision detection method to dynamically step time and guarantee intersection-free simulation for elastodynamics contact problems with friction across variable material parameters and mesh resolutions. We hypothesized IPC’s ability to robustly simulate contact-rich dynamics could be used to run robust dynamic grasp simulations, and generate more accurate predictions of grasp outcomes observed on a physical system.

We adapt IPC to the problem of Sim2Real robot grasping of rigid objects with soft gripper pads to create IPC-GraspSim. IPC-GraspSim modifies IPC’s state machine by introducing a piecewise grasping action that 1) dynamically adapts the grasp width to the local object geometry by iteratively closing a set of gripper jaws until a deformation energy threshold is reached due to contact with the object, 2) lifts the jaws to a specified height, and 3) records grasp success based on the relative final positions of the jaws and object. IPC-GraspSim

accurately models grasp outcomes produced by a physical system, as shown in Fig. 1. This thesis makes the following contributions:

1. IPC-GraspSim: an application of the contact dynamics model introduced in Incremental Potential Contact [27] to the Sim2Real robot grasping problem that models both friction and contact dynamics effects for compliant parallel-jaw gripper jaws.
2. An Isaac Gym-based parallel-jaw grasping simulator with FleX backend that extends Eppner et al. [10], which only supports rigid gripper - rigid object interactions, by adding deformable gripper jaws.
3. Experiments comparing IPC-GraspSim with Isaac Gym and 3 analytic contact models (soft point contact, REACH, and 6DFC), that suggest IPC-GraspSim can predict grasp robustness with F1 score of 0.85, an improvement in F1 score by 0.03 to 0.20 over analytic baselines and 0.09 over Isaac Gym, at a cost of 8000x more compute time.

All data, code, videos, and supplementary material are available at <https://sites.google.com/berkeley.edu/ipcgraspsim>.

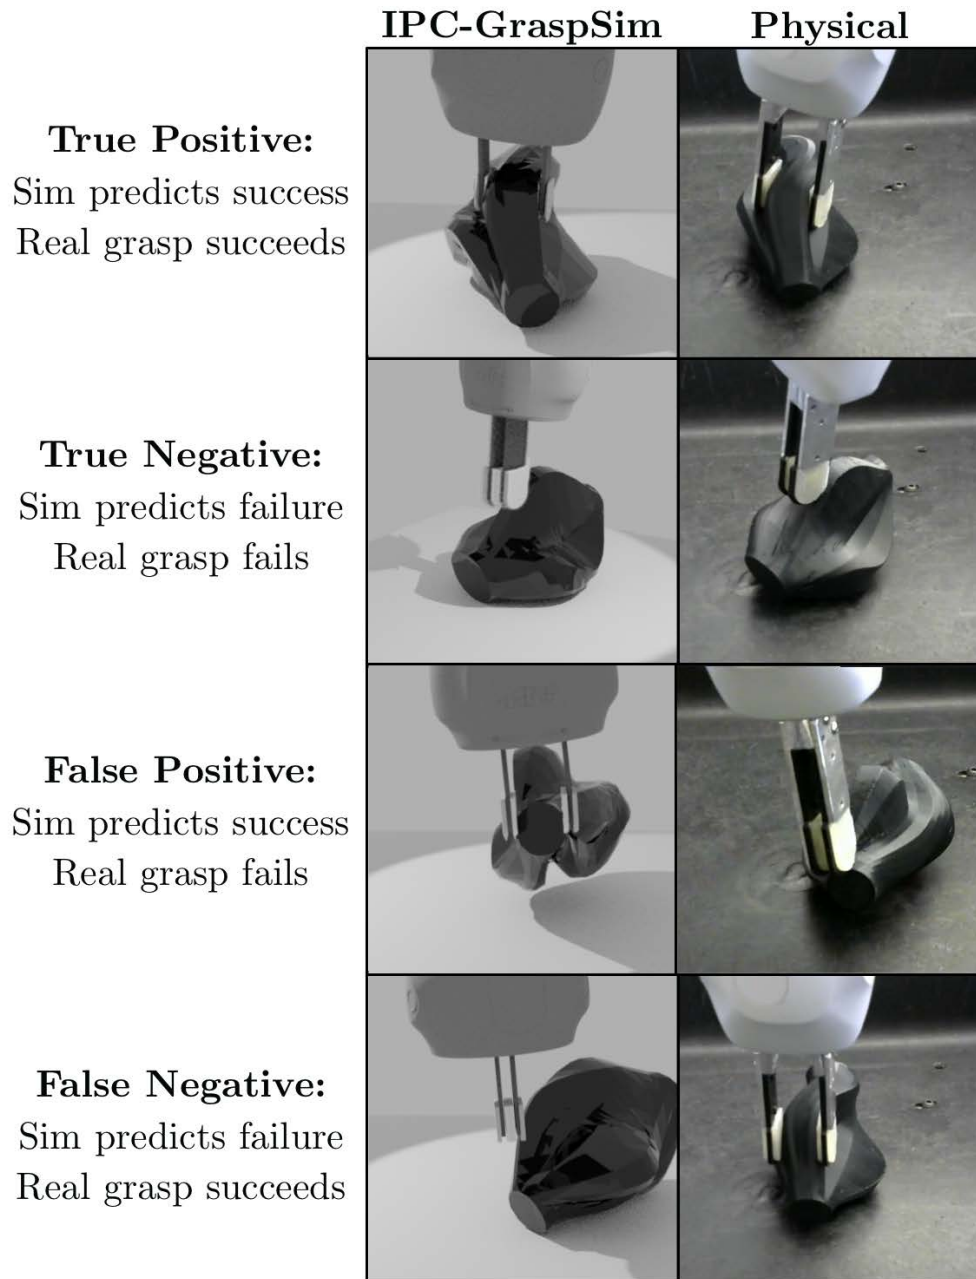


Figure 1.1: Simulated (left) and physical (right) grasp outcomes: in the first two rows, IPC-GraspSim correctly predicts grasp success and grasp failure. In the last two rows, predictions are incorrect due to unmodeled cantilever bending + rotations due to pushing.

Chapter 2

Related Work

2.1 Grasp Metrics and Contact Models for Soft Grippers

Analytic grasp metrics use contact force and torque constraints generated by contact models to measure grasp robustness. As summarized by Roa et al. [40], common metrics describe the ability of a grasp to resist external disturbances (force closure [3, 39]), be at equilibrium between contact forces and external forces [21], or allow the object to resist gravity (wrench resistance [32, 33]); these metrics can be determined using linear or quadratic programs based on the force and torque constraints at each contact.

Grasp contact models measure the ability of grasp contacts between the gripper jaw and the object to resist external forces and torques via the contact area, contact pressure and friction applied at the contact. Models of compliant gripper contacts typically use the quasistatic approximation. Bicchi et al. [4] summarize analytic physics-based contact models, including soft point contact models [6, 15, 19, 20, 43], which exert forces in the plane tangent to the contact surface and a torsional moment about the contact normal. Frictional limit surface models jointly limit the tangential force and torque about the contact normal [12, 15, 28, 49] and can be extended to model non-planar contacts [46]. Another approach is to approximate contact area with geometric primitives, such as planes, spheres, or cylinders, that have analytic solutions [5, 14]. Similarly, Danielczuk et al. [8] and Xu et al. [45] discretize the contact area into a triangular mesh and formulate per-triangle or per-contact limit surface constraints based on a static friction assumption. In contrast to these quasistatic grasp contact models and metrics, we use a simulator that incorporates both grasp dynamics and frictional area contacts.

2.2 Grasp Simulators for Soft Grippers

Grasp simulators, unlike grasp contact models, can incorporate dynamics which are crucial to grasping: they can predict which jaw makes contact first, and how that affects post-grasp displacement as the gripper and object settle into a stable position [48]. In this paper, we focus on dynamic simulators that model deformable materials [7, 26, 29, 35], as it is important to model frictional contact dynamics simultaneously with gripper compliance: deformation affects surface area, which in turn affects object dynamics.

Isaac Gym is a GPU-enabled simulator favored by the robotics community for its speed and python-based simulator with a straightforward robotics interface. NVIDIA FleX engine integration has recently made deformables available for robotics environments, as evidenced by recent papers that model the SynTouch BioTac [37] and grasping of deformable objects with rigid grippers [16]. However, many simulators, including Isaac Gym, require manual parameter tuning to avoid interpenetration artifacts and may be sensitive to friction and closing force parameters [41].

In simulation, grasp success may be defined as (1) the object can be lifted above the surface [16, 31], (2) the object can be lifted and withstand perturbations via shaking motions [9, 10], (3) the object is removed from its environment after grasp and move actions [47], or (4) the object’s post-grasp pose matches a desired pose [1]. In this paper, we use the first definition of grasp success.

2.3 Simulation of Deformable Materials

Simulating deformable materials is challenging due to their many degrees of freedom and friction between deformable and rigid objects. However, due to the high computational load in more accurate representations such as Finite Element Methods (FEM), most work relies on simplified models as provided by mass-spring systems and position-based dynamics, which are computationally efficient but do not have a rigorous mathematical model for contact and friction [2]. With FEM (e.g., PyBullet’s FEM option and graphics implementations), Neo-Hookean materials are often used to simulate hyperelastic materials where the stress-strain relationships are highly nonlinear, such as silicone rubber [34]. Isaac Gym models softbody dynamics using corotational linear-elastic constitutive models, and leverages a custom GPU-based solver [30].

A common model for applying loading conditions onto the deformable bodies is the penalty-based contact, which calculates contact force as proportional to the mesh penetration depth and penalty value p , assuming nonzero interpenetration for all contacting bodies. High values of p can decrease the penetration depths, but also make the problem ill-conditioned [22]. Furthermore, any contact violation can cause catastrophic failures [27]. A common way to handle interpenetration is to add a boundary layer on the deformable to “thicken” the mesh for collisions, but the layer depth must be manually tuned for each scenario [29, 41].




Figure 2.1: Cube grasps with incorrect contact handling: (A) interpenetration, (B) force-field effect.

Contact Handling for Grasping

The computationally efficient penalty-based contact methods are prone to large interpenetration effects, as shown in Figure 2.1A: the gripper pad intersects with the target object. These effects become more extreme when simulating softbody contacts with large Young’s Modulus values E , or contacts between materials with large differences in E .

In order to avoid large interpenetration errors, other works have modified penalty-based contact methods to calculate contact forces based on a “thickened” version of the mesh, such that nonzero contact forces can be applied without observing any mesh penetration. However, this method may introduce “force field” effects if the layer depth is too large: an example is shown in Figure 2.1B, where the gripper pad is able to push and grab onto the target object although no contact is visually observed.

Neither of these behaviors are not acceptable in grasping contexts, as the error in Figure 2.1A will lead to false sticking behavior (false positive), and Figure 2.1B will lead to incorrect dynamics from inaccurate collision handling.

2.4 Incremental Potential Contact

Incremental Potential Contact (IPC) [27] is a computer graphics model and algorithm presented in 2020 based on a custom nonlinear Newton-based solver that solves elastodynamics contact problems with the smoothed barrier method. It guarantees a collision-free state at every timestep to provide robust solutions for all choices of material parameters, timestep sizes, impact velocities, deformation severities, and enforced boundary conditions. IPC provides inversion-free and intersection-free solutions, so that objects do not penetrate and




Figure 2.2: “Squishy ball test” figure from IPC [27]: a soft toy with many tendrils is thrown into a glass wall with high velocity. The custom collision detection implementation and contact modeling allows the microstructures to disentangle themselves post-impact with no interpenetration artifacts.

remain stuck to one another in future timesteps. IPC accounts for static friction with a smoothed approximation to eliminate an explicit Coulomb constraint, which leads to the assumption that static and dynamic friction coefficients are equal between each material combination. The authors justify this decision by implementing an approximation to static friction that can successfully model stiction. This implementation allows the model to robustly simulate scenes that require softbody-softbody collisions with high contact forces, as shown in Figure 2.2

We hypothesized that we could more accurately model grasp outcomes observed on a physical system by leveraging IPC’s intersection-free guarantee, as the model can robustly model contact-heavy dynamics that occur during robot grasping while avoiding interpenetration issues described in Section 2.3. We test this hypothesis by implementing a grasping simulator that uses IPC as the physics engine, and benchmark its performance with a physical grasp dataset collected by Danielczuk et al. [8].

Chapter 3

IPC-GraspSim

3.1 Problem Statement

We evaluate the ability of IPC-GraspSim to accurately predict the grasp outcomes in a physical grasp dataset with known object and gripper geometries and known grasp poses.

Assumptions

We make the following assumptions:

1. A singulated rigid object of known geometry and uniform density rests on a planar surface.
2. The gripper has known geometry and two parallel jaws.
3. The silicone rubber material on the gripper jaws can be simulated with the Neo-Hookean material model.

Definitions

The true state \mathbf{x} of the grasping environment includes geometry (including deformations), pose, mass, and frictional properties of the object and gripper jaws. Each grasp \mathbf{u} is represented by the parallel gripper jaw center and the grasp axis orientation $\varphi \in \mathbb{R}^3$. The modifiable gripper jaw parameters include the deformable material elasticity and the friction coefficient between the gripper jaw pads and the grasping object. Success $S(\mathbf{x}, \mathbf{u})$ is measured with a binary reward function S , where $S = 1$ if the grasp successfully lifts the object and $S = 0$ otherwise.

To model uncertainty and imprecision in robot control, we define grasp robustness $R(\mathbf{x}, \mathbf{u}) = \mathbb{E}[S(\mathbf{x}, \mathbf{u}')]$, where the perturbed grasp pose $\mathbf{u}' = \mathbf{u} + \varepsilon$ is offset by ε sampled from a normally-distributed pose random variable. To calculate robustness in our experiments, we use Monte-Carlo sampling to estimate R with the sample mean of N trials each with a perturbed gripper




Figure 3.1: Two grasps on a cube with IPC-GraspSim, showing state at successive timesteps. (Top row) Successful grasp: 1) cube and gripper initialized; 2) cube contacts right jaw; 3) cube contacts left jaw; 4) gripper lifts cube. (Bottom row) Failed grasp: 1) cube and gripper initialized; 2) cube contacts right jaw; 3) cube *misses* left jaw; 4) jaws meet each other.

pose: $R = \frac{1}{N} \sum_{i=1}^N S(\mathbf{x}, \mathbf{u}'_i)$. We emphasize the importance of using sampling with grasp trials; although the robot may intend the same grasp on a given object, not all trials will behave equally due to randomness introduced by pose estimation, robot control, and frictional forces.

Metrics

To quantify how closely the predicted grasp robustness from the simulator matches the robustness measured on the physical system, we measure the F1 score, a harmonic mean of average precision (AP) and average recall (AR). F1 score is a common measure of binary classification performance for imbalanced datasets in computer vision [42].

3.2 IPC-GraspSim

To simulate grasps, we apply Incremental Potential Contact [27], as it guarantees intersection-free and inversion-free solutions for all parameter choices, as described in Section 2.3. This model avoids the interpenetration artifacts faced by many simulator models [7, 29], which include false sticking behavior after gripper lifting (false positive) or intersecting through the object (false negative). To avoid these large interpenetration errors in other simulators,




Figure 3.2: Corner grasp on cube: (left) gripper perturbed away from object, physical grasp attempt fails (right) gripper perturbed towards the object, physical grasp attempt succeeds. Both grasp trials intend to perform the same action, yet uncertainties allow the grasp success to vary across trials.

collision parameters need to be carefully tuned. As IPC is originally created for computer graphics applications, IPC-GraspSim implements a robotic grasp interface: the inputs for IPC-GraspSim are the grasping state \mathbf{x} and grasp \mathbf{u} . Then the simulator performs a perturbed grasp $\mathbf{u}' = \mathbf{u} + \varepsilon$ where ε is noise in grasp pose, with translation noise (in meters) sampled from $\mathcal{N}(0, 0.001)$ and rotation noise (in radians) from $\mathcal{N}(0, 0.003)$. IPC-GraspSim returns grasp evaluation $S(\mathbf{x}, \mathbf{u}')$ upon the end of the simulation, and we retrieve robustness $R(\mathbf{x}, \mathbf{u})$ over multiple simulation trials with different perturbations. We also implement a finite state machine in the original IPC simulator in order to progress through a piece-wise grasping motion that actively adapts to the grasp state \mathbf{x} , which is detailed in the following section. No changes are made to IPC’s contact dynamics model.

Stochasticity in Sim and Real

It is critical for grasp simulators to account for uncertainty and imprecision in robot control, as there are two main sources of uncertainties when grasping in physical environments: robot position and object pose estimation. For some grasps that are very sensitive to contact area, as shown in Figure 3.2, perturbations have higher chances of affecting grasp success rates.

The calculated post-grasp object pose and grasp success in IPC-GraspSim is deterministic for each input grasp. To introduce randomness for Monte-Carlo sampling in this simulator, we inject grasp pose perturbation matrix ε with normally distributed translation and rotation error. Although Isaac Gym (with FleX backend) is not deterministic, we inject the same perturbation effects in the simulator to mirror our changes in IPC-GraspSim.

Grasp Simulation Procedure

IPC-GraspSim models grasping with a velocity-controlled compliant parallel-jaw gripper as follows:

1. Initialize

Initialize gripper jaws and object location as specified by the perturbed grasp \mathbf{u}' and state \mathbf{x} . Gripper jaws are opened to their max grasp width, and unobstructed by any object other than the target. The gripper is modeled as two identical jaws, each composed of an extruded compliant polygon pad attached to a rigid rectangular plate. Figures 1 and 3.1 show the compliant padding in light gray, and rigid backing in dark gray.

2. Squeeze

Close gripper pads along the grasp axis with a constant velocity. The target object will interact with the gripper pads and the ground surface, and is free to translate or rotate.

Although YuMi, the physical robot used in this paper, completes the gripper closing action when it reaches a maximum specified closing force, contact force information is currently unavailable in the IPC simulator. Due to the limitations in the IPC simulator, IPC-GraspSim uses Ψ , the deformation or strain energy in a given mesh, to approximate closing force applied by the compliant gripper jaws on the target object.

IPC-GraspSim approximates closing force on the target object by checking if A) the deformation energy Ψ described in Li et al. [27] in both gripper tooltips Ψ_1, Ψ_2 have increased such that $\Psi_1, \Psi_2 \geq \Psi_{th}$, and B) the energy is similar in both gripper jaws $\Psi_1 \approx \Psi_2$. Ψ_{th} is a user-defined threshold to limit the amount of deformation in each pad (Ψ_1, Ψ_2) and corresponds to the Ψ that best matches maximum gripper closing force for the specified gripper. (A) assures that the closing action will not incorrectly terminate due to deformations driven by other external forces (e.g., gravitational), and (B) assures that equal effort is applied by both grippers. We test various values of Ψ_{th} on a range of grasps. Based on visual comparison with our physical silicone gripper pads, we choose $\Psi_{th} = 5 \times 10^4$.

We attempt multiple grasps on the ‘cube’ object with various values of Ψ_{th} (shown in Figure 3.3). We find that $\Psi_{th} = 5 \times 10^4$ properly models the gripper pad thinning under pressure while avoiding unrealistic deformation.

3. Lift

Once the deformation energy threshold is reached, the two gripper jaws are lifted upwards with the jaw distance held constant until the object’s lowest point is at least 2 cm above the plane by the time the simulation finishes.

A simulated grasp is considered successful if both gripper jaws are in contact with the object at the end of simulation, and the object is not in contact with the plane. Examples of successful and failed grasps are shown in Figure 3.1.




Figure 3.3: Tested values for Ψ_{th} , and the gripper pad deformation. Sides of the gripper pad are colored in red. $\Psi_{th} = 5 \times 10^4$ models the gripper pad thinning under pressure, without deforming in an unrealistic way.

3.3 Parameter Sensitivity Analysis

We analyze the sensitivity of IPC-GraspSim to the elasticity E and dynamic friction μ of the gripper pad material using a coarse grid search on the parameter space. We also test two grippers with different geometries, as shown in Table 3.1 (rectangular) and Table 3.2 (rounded), to study the effect of accurately modeling the gripper geometry.

Test Object Set

In our experiments, we use the physical dataset from Danielczuk et al. [8]: a set of 2,000 grasps across a set of 16 3D-printed adversarial objects on a physical ABB YuMi robot with a compliant parallel-jaw gripper [13], shown in Figure 3.4.

This dataset contains 21 adversarial objects where analytic grasp methods perform poorly due to large post-grasp displacements. Even for seemingly simple objects, such as object 5 (“cube”) and object 10 in Figure 3.4, small perturbations as mentioned in Section 3.2 introduce discrepancies between pre-grasp and post-grasp contact area. However, a subset of 16 objects were chosen for our experiments due to 1) skewed success/unsuccessful grasp ratio and 2) timing constraints, as IPC-GraspSim was not optimized for speed.

We evaluate grasp outcomes on 3 objects with 1) a similar number of successful and




Figure 3.4: The 21 adversarial objects used in the grasp dataset collected by Danielczuk et al. [8], where analytic grasp models perform poorly. Each object is sampled with 25 test grasps, 5 trials per grasp. The training object set (cat, pear, vase) for our experiments is highlighted in purple.

unsuccessful grasps and 2) varied geometries. Then, we identify the set of parameters $(\hat{E}, \hat{\mu})$ with the highest F1 score in the training set, and use these parameters to simulate grasps on the remaining 13 objects in Section 4.

Sensitivity to Young’s Modulus and Friction


We vary the Young’s modulus of the jaw pad from 10^7 Pa (rubbery material) to 10^{11} Pa (metal material). We linearly vary the friction coefficient between the pad and object from 0.3 to 0.6. Note that the exact friction between the jaw pad and object is unknown due to unknown material properties and the unmodeled surface variations on the jaw pad, as shown in Dex-Net’s silicone rubber fingers. The millistructure design results in increased friction between the jaw pad and the object [13]. We aim to find both the Young’s modulus and single friction coefficient that closely models such a structure.

For both gripper jaw models, lower Young’s modulus values $(E, \mu) = (10^8, 0.4)$ have higher AP, AR, and F1 score as compared to higher Young’s modulus values $(10^{11}, 0.4)$. For rounded jaws, more compliant jaws increase AP by 13%, AR by 2%, and F1 by 7%; for rectangular jaws, they increase AP by 14%, AR by 11%, and F1 by 12%. Increase in μ appears to correlate with an increase in AR, as the additional frictional forces allow for a greater number of successful grasps; however, AP drops as the number of false positives increases. We hypothesize the increase in AR occurs due to the increased contact area in compliant grippers, allowing the grippers to deform around the contact point and distribute contact forces across the object’s surface. However, when μ is too high, objects no longer can slip out of the gripper jaws and the number of false positive predictions increases.



	Young's modulus E (Pa)	Friction coefficient μ			
		0.3	0.4	0.5	0.6
1e7		0.77	0.82	0.83	0.84
1e8		0.85	0.86	0.80	0.78
1e9		0.84	0.81	0.77	0.74
1e10		0.77	0.78	0.75	0.74
1e11		0.77	0.74	0.73	0.74

Table 3.1: **Rectangular** jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 0.86.



	Young's modulus E (Pa)	Friction coefficient μ			
		0.3	0.4	0.5	0.6
1e7		0.66	0.70	0.70	0.74
1e8		0.85	0.91	0.89	0.85
1e9		0.85	0.82	0.86	0.83
1e10		0.82	0.88	0.87	0.84
1e11		0.82	0.84	0.81	0.86

Table 3.2: **Rounded** jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 0.91.

Gripper Geometry

The compliant rounded jaws result in a 17% increase in AP and 8% decrease in AR, leading to an overall increase of 5% in F1 score, compared to the compliant rectangular jaws when using the highest-performing parameters for each geometry ($E = 10^8, \mu = 0.4$). We hypothesize that the slight precision-recall tradeoff occurs as the smaller cross-sectional area of the rounded gripper jaw causes grasp robustness predictions to be more conservative. However, the F1 score ultimately increases, as the extra contact area due to the rectangular jaws had allowed grasps to incorrectly succeed.

Based on the parameter sensitivity analysis, we use the rounded gripper jaw model with $(\hat{E}, \hat{\mu}) = (10^8, 0.4)$ to model the compliant gripper jaws in dynamic simulation. We note that this optimal set of parameters closely matches the physically estimated parameters ($10^7 - 10^8$ Pa, 0.46).

Chapter 4

Experiments

4.1 Experiment Setup

To evaluate the precision and recall of IPC-GraspSim, we use the 13 held-out objects from physical dataset from Danielczuk et al. [8]. The objects in this dataset were chosen because they span a wide range of object geometries and grasp success rates. All objects are 3D printed such that simulated and physical attempted grasps are identical. For benchmark experiments, we run 5 trials for each grasp while perturbing the grasp pose in each trial for dynamics randomization. The trials are then averaged to calculate the grasp robustness predicted by each grasp model.

4.2 Benchmark Grasp Models

In addition to IPC-GraspSim, we benchmark 4 other models; three analytic models used in the literature (soft point contact model, REACH [8], 6DFC [45]), and a dynamic softbody simulation implemented in NVIDIA Isaac Gym with Flex backend [29]. We measure the performance of Isaac Gym and IPC-GraspSim by simulating grasps with the rounded gripper jaws and material properties ($E = 10^8, \mu = 0.4$) determined in Section 3.3. Soft Point, REACH, 6DFC, and IPC-GraspSim are run on an Ubuntu 18.04 on a single core of a Dual 20-Core Intel Xeon E5-2698 v4 2.2 GHz, and Isaac Gym on an Ubuntu 20.04 machine with an Intel i7-6850K processor and NVIDIA Titan X GPU.

The results in Table 4.2 suggest IPC-GraspSim can predict grasp robustness with F1 score of 0.85, an improvement in F1 score by 0.03 to 0.20 over analytic baselines and 0.09 over Isaac Gym, at a cost of 8000x more compute time. We observe an average increase of 0.17 in average recall in IPC-GraspSim compared to grasp models that approximate compliance (Soft Point, REACH, 6DFC) and a 0.05 increase in average recall compared to Isaac Gym.

We hypothesize that the increase in recall as compared to the analytic models is due to the ability of IPC-GraspSim to model dynamics and compliance simultaneously. Modeling dynamics results in correct predictions for grasps that rotate into alignment with the gripper,

Grasp config (5 trials/grasp)	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
IPC-GraspSim	0	0	1	0	0	1	0.4	0.8	1	1	1	1	1	1	0.6	0	0	0	0.4	0	0	1	1	0	0
Physical Experiments	0	1	1	0	0	1	0.8	0	1	1	1	1	1	1	0	0	0.2	0	0	0.4	1	0.8	0	0	

Table 4.1: Grasp robustness values predicted by IPC-GraspSim (top row) and observed from 125 physical experiments (bottom row) using an ABB YuMi with a compliant parallel-jaw gripper on the “vase” object shown in Fig. 1. Each column corresponds to a unique grasp configuration (i.e., relative pose between gripper and object). Each cell contains the robustness $R = \frac{1}{N} \sum_{i \in \{1..N\}} S_i$, with $N = 5$, and is colored to indicate its relative robustness (red: unreliable, green: robust). We observe a very high correlation of predicted grasp robustness between the two rows. Example true positive (grasp 2), true negative (grasp 0), false positive (grasp 7), and false negative (grasp 1) grasps are shown in Fig. 1. See Section 4.3 for discussion of grasp configurations 1 and 7.

Model	AP	AR	F1	Runtime (s)
Soft Point	0.89	0.59	0.71	0.02
REACH [8]	0.75	0.57	0.65	0.44
6DFC [45]	0.82	0.83	0.82	0.07
Isaac Gym	0.75	0.79	0.76	421
IPC-GraspSim	0.86	0.84	0.85	605

Table 4.2: Average Precision (AP) and Average Recall (AR) for each model’s grasp quality predictions for the 2000 grasps collected on the physical robot. All runtimes are provided per-grasp. Per-frame runtimes for Isaac Gym and IPC-GraspSim are 0.28 and 6.05 seconds respectively.

a common failure mode for quasistatic models [8, 45]. In addition, modeling compliance keeps the object closer to the gripper as it moves by 1) conforming to the shape of the object and 2) applying normal and frictional forces that stabilize the object. The combination of dynamics and compliance models the changing contact area and contact pressure over the course of the grasping motion.

We also note that despite modeling dynamics, Isaac Gym has F1 score 0.06 lower than 6DFC, a quasistatic contact model; this result suggests that if the underlying simulator for modeling compliant gripper grasps is not sufficiently robust (i.e., simulator returns sudden large forces due to incorrect collision force handling), it may perform worse than a robust model that ignores dynamics. Although IPC-GraspSim only increases F1 score by 0.03 compared to 6DFC [45], the simulator has significant room for improvement: most of its failure modes are due to unmodeled physical effects (e.g., gripper jaw cantilever effects, further detailed in Section 4.3), which we hope to address in future work.

In comparison to Isaac Gym, IPC-GraspSim increases average precision by 0.11, average

recall by 0.05, and F1 score by 0.09. The performance of IPC-GraspSim highlights the importance of intersection-free and robust modeling of soft contacts for compliant gripper modeling. First, poorly-resolved contacts may cause grasped objects to move unpredictably (e.g., float or swing into air, as shown in Figure 4.1C-D for Isaac Gym). Second, incorrect contacts introduce force imbalance between the two compliant jaws and cause the target object to wobble around or become lodged in the gripper, resulting in unrealistic object behavior.

Although dynamic simulators (Isaac Gym, IPC-GraspSim) run at least three orders of magnitudes slower than analytic grasp models, these evaluations are intended to run offline for building large-scale datasets. Therefore, these simulations can be performed in parallel over extended periods of time.

4.3 Failure Mode Analysis

Table 4.1 shows results for each grasp attempt of the “vase” (object 21 in the grasp dataset collected by Danielczuk et al. [8], shown in Figure 1). A full breakdown of grasp attempts and predictions for all objects can be found in the supplement. For this object, we observe a high correlation between IPC-GraspSim predictions and physical experiments. IPC-GraspSim achieves an F1 score of 0.85 as compared to 0.65-0.82 over analytic methods and 0.68 for Isaac Gym.

We observe a false positive prediction for grasp 7, where gripper jaws bend and cantilever away from the object under the stress from the closing forces. The contact forces then push the vase away from and out of the gripper. However, in simulation, the gripper backings are perfectly rigid, and do not account this bending phenomenon.

We also observe a false negative prediction for grasp 1, where the right jaw contacts the vase first and rotates the vase out of the gripper. This phenomenon is also observed in Figure 4.1A. This missed prediction suggests that precise gripper geometry and gripper closing speed have an impact on grasp predictions. We additionally hypothesize that the dynamics between the ground plane and the object are not well-modeled in simulation.




Figure 4.1: Four grasps (columns) with three simulators: (Top to bottom) Isaac Gym, IPC-GraspSim, Physical Experiments. (A) False negative: IPC-GraspSim and Isaac Gym fail to predict that the right jaw will contact the cube. (B) True positive: bowl correctly moves and rotates into gripper for both sims. (C) False positive: In Isaac Gym, incorrect contact forces push “pawn” up and away from the plane. In IPC-GraspSim, rigid jaws have sufficient contact area to grasp the object. Physical grasp fails as jaw cantilevering decreases contact area. (D) False negative: For Isaac Gym, the object is pushed upwards, with a similar issue as described in (C). In IPC-GraspSim, left gripper jaw pushes object out of the gripper. Videos of all simulator grasp attempts can be found on the website.

Chapter 5

Discussion

5.1 Conclusion

In this paper, we present IPC-GraspSim, an intersection-free, robust grasp simulator, and evaluate its performance on a physical grasp dataset generated with compliant parallel-jaw grippers. We make the insight that Incremental Potential Contact’s unique property — ability to robustly simulate large-deformation dynamics without intersections — could be used to simulate contact-rich dynamics during robot grasping with compliant gripper jaws. Compared with analytic baselines, IPC-GraspSim improves recall while maintaining similar precision, which we hypothesize is a result of modeling the compliance in a dynamic context: the contact points between the grippers and the target object change throughout the entire gripper closing process. IPC-GraspSim also improves precision over Isaac Gym due to its ability to stably model soft contacts without interpenetration. IPC is designed for simulating contact dynamics, and thus puts speed as lower priority.

5.2 Future Work

As future work, we started a collaboration with authors from IPC and Medial-IPC [23] to investigate an extension of IPC-GraspSim with Affine Body Dynamics (ABD) by Lan et al. [24]. This grasp simulator requires significantly less compute time to label grasps for objects of varying complexities, demonstrating speed performance improves up to 100x. We hope that the lower computational cost of this simulator will also allow us to explore further into IPC-GraspSim failure cases due to model mismatch.

We will also measure the simulator’s performance for grasps with different gripper and object geometries, to study how well this simulator extends to unseen robot grasping environments. We will also explore the effect of training grasp prediction neural networks using grasp labels generated with this new simulator; new grasp sampling methods (i.e., not antipodal) may be required to leverage in-hand object dynamics during grasping with gripper jaw pushes.

5.3 Reflections

In completing this thesis, I experienced first-hand both the strengths and limitations of simulations for robotics research: although there is a large support for computationally efficient articulated rigid-body systems (e.g., robot arms), softbody simulations remain brittle and unreliable. I also learned how different fields prioritize different aspects of physics-based simulations: simulations for robotics (e.g., Isaac Gym) focus on computational efficiency for rigid-body systems, at a cost of physical accuracy; simulations for engineering focus on physical accuracy; and simulations for computer graphics focus on visual accuracy. It still amazes me that a computer graphics simulator can be used to label grasp robustness, a robot task that demands highly accurate physic-based modeling.

This last year in AUTOLab helped me grow as a researcher and collaborator, and prepare to tackle a PhD. From Professor Goldberg, I learned the importance of presenting your work clearly. By refining the motivations and descriptions for research projects, I feel that my work is understood much better by not just others, but also by myself. From Mike, I learned the importance of software engineering to conduct research, and making modular codebases to reuse and accelerate future research directions.

Finally, from the research community at UC Berkeley, I learned the importance of “talking it out”: during the course of my masters program, I frequently wondered — what could have I done differently to do better research? I found the answer to this question by spending time with my research lab colleagues, from pulling all-nighters together through conference deadlines and skiing together during impromptu Tahoe trips. From these professional and casual conversations over the last few months, I learned that great research comes from not just focusing on the problem at hand, but by also taking the time to communicate my ideas to others.

Fortunately, I will have the opportunity to apply these lessons to my future research during my PhD at Berkeley over the next few years.

Bibliography

- [1] Arthur Allshire et al. “Transferring Dexterous Manipulation from GPU Simulation to a Remote Real-World TriFinger”. In: *Conf. on Robot Learning (CoRL)*. 2021.
- [2] Jan Bender et al. “A survey on position-based simulation methods in computer graphics”. In: *Computer graphics forum*. Vol. 33. 6. Wiley Online Library. 2014, pp. 228–251.
- [3] Antonio Bicchi. “On the closure properties of robotic grasping”. In: *Int. Journal of Robotics Research (IJRR)* 14.4 (1995), pp. 319–334.
- [4] Antonio Bicchi and Vijay Kumar. “Robotic grasping and contact: A review”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. Vol. 1. 2000, pp. 348–353.
- [5] Matei Ciocarlie, Claire Lackner, and Peter Allen. “Soft finger model with adaptive contact geometry for grasping and manipulation tasks”. In: *Proc. IEEE Eurohaptics Conf. and S. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 2007, pp. 219–224.
- [6] Matei Ciocarlie, Andrew Miller, and Peter Allen. “Grasp analysis using deformable fingers”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2005, pp. 4122–4128.
- [7] Erwin Coumans and Yunfei Bai. *PyBullet, a Python module for physics simulation for games, robotics and machine learning*. <http://pybullet.org>. 2016–2021.
- [8] Michael Danielczuk et al. “Reach: Reducing false negatives in robot grasp planning with a robust efficient area contact hypothesis model”. In: *Int. S. Robotics Research (ISRR)*. 2019.
- [9] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. “A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set”. In: *Int. S. Robotics Research (ISRR)*. 2019.
- [10] Clemens Eppner, Arsalan Mousavian, and Dieter Fox. “ACRONYM: A Large-Scale Grasp Dataset Based on Simulation”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2020.
- [11] Hao-Shu Fang et al. “Graspnet-1billion: A large-scale benchmark for general object grasping”. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. 2020, pp. 11444–11453.

- [12] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. “Planar sliding with dry friction part 1. limit surface and moment function”. In: *Wear* 143.2 (1991), pp. 307–330.
- [13] Menglong Guo et al. “Design of parallel-jaw gripper tip surfaces for robust grasping”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2017, pp. 2831–2838.
- [14] Heinrich Hertz. *Miscellaneous papers*. Macmillan, 1896.
- [15] Robert D Howe, Imin Kao, and Mark R Cutkosky. “The sliding of robot fingers under combined torsion and shear loading”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 1988, pp. 103–105.
- [16] Isabella Huang et al. “DefGraspSim: Simulation-based grasping of 3D deformable objects”. In: *DO-Sim: Workshop on Deformable Object Simulation in Robotics at Robotics: Science and Systems (RSS)*. 2021.
- [17] Edward Johns, Stefan Leutenegger, and Andrew J Davison. “Deep learning a grasp function for grasping under gripper pose uncertainty”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2016, pp. 4461–4468.
- [18] Dmitry Kalashnikov et al. “Scalable deep reinforcement learning for vision-based robotic manipulation”. In: *Conf. on Robot Learning (CoRL)*. 2018, pp. 651–673.
- [19] Imin Kao and Mark R Cutkosky. “Quasistatic manipulation with compliance and sliding”. In: *Int. Journal of Robotics Research (IJRR)* 11.1 (1992), pp. 20–40.
- [20] Imin Kao and Fuqian Yang. “Stiffness and contact mechanics for soft fingers in grasping and manipulation”. In: *IEEE Trans. Robotics and Automation* 20.1 (2004), pp. 132–135.
- [21] Jeffrey Kerr and Bernard Roth. “Analysis of multifingered hands”. In: *Int. Journal of Robotics Research (IJRR)* 4.4 (1986), pp. 3–17.
- [22] G. Kloosterman. “Contact methods in finite element simulations”. Undefined. PhD thesis. Dec. 2002.
- [23] Lei Lan et al. “Medial IPC: accelerated incremental potential contact with medial elastics”. In: *ACM Transactions on Graphics (TOG)* 40.4 (2021), pp. 1–16.
- [24] Lei Lan et al. “Affine Body Dynamics: Fast, Stable & Intersection-free Simulation of Stiff Materials”. In: *arXiv preprint arXiv:2201.10022* (2022).
- [25] Ian Lenz, Honglak Lee, and Ashutosh Saxena. “Deep learning for detecting robotic grasps”. In: *Int. Journal of Robotics Research (IJRR)* 34.4-5 (2015), pp. 705–724.
- [26] Beatriz León et al. “Opengrasp: a toolkit for robot grasping simulation”. In: *Proc. IEEE Int. Conf. on Simulation, Modeling, and Programming of Autonomous Robots (SIMPAN)*. Springer. 2010, pp. 109–120.
- [27] Minchen Li et al. “Incremental potential contact: Intersection-and inversion-free, large-deformation dynamics”. In: *ACM transactions on graphics* (2020).

- [28] Yanmei Li and Imin Kao. “A review of modeling of soft-contact fingers and stiffness control for dextrous manipulation in robotics”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. Vol. 3. 2001, pp. 3055–3060.
- [29] Jacky Liang et al. “GPU-accelerated robotic simulation for distributed reinforcement learning”. In: *Conf. on Robot Learning (CoRL)*. 2018, pp. 270–282.
- [30] Miles Macklin et al. “Non-smooth newton methods for deformable multi-body dynamics”. In: *ACM Transactions on Graphics (TOG)* 38.5 (2019), pp. 1–20.
- [31] Jeffrey Mahler and Ken Goldberg. “Learning deep policies for robot bin picking by simulating robust grasping sequences”. In: *Conf. on Robot Learning (CoRL)*. 2017, pp. 515–524.
- [32] Jeffrey Mahler et al. “Dex-Net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2018.
- [33] Jeffrey Mahler et al. “Learning ambidextrous robot grasping policies”. In: *Science Robotics* 4.26 (2019), eaau4984.
- [34] Gilles Marckmann and Erwan Verron. “Comparison of hyperelastic models for rubber-like materials”. In: *Rubber chemistry and technology* 79.5 (2006), pp. 835–858.
- [35] Andrew T Miller and Peter K Allen. “Graspit! a versatile simulator for robotic grasping”. In: *IEEE Robotics & Automation Magazine* 11.4 (2004), pp. 110–122.
- [36] Douglas Morrison, Peter Corke, and Jürgen Leitner. “Learning robust, real-time, reactive robotic grasping”. In: *Int. Journal of Robotics Research (IJRR)* 39.2-3 (2020), pp. 183–201.
- [37] Yashraj Narang et al. “Sim-to-real for robotic tactile sensing via physics-based simulation and learned latent projections”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2021.
- [38] Lerrel Pinto and Abhinav Gupta. “Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2016, pp. 3406–3413.
- [39] Elon Rimon and Joel Burdick. “On force and form closure for multiple finger grasps”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. Vol. 2. 1996, pp. 1795–1800.
- [40] Máximo A Roa and Raúl Suárez. “Grasp quality measures: review and performance”. In: *Autonomous robots* 38.1 (2015), pp. 65–88.
- [41] Alessio Rocchi et al. “Stable simulation of underactuated compliant hands”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2016, pp. 4938–4944.
- [42] Takaya Saito and Marc Rehmsmeier. “The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets”. In: *PloS one* 10.3 (2015), e0118432.

- [43] J Kenneth Salisbury and B Roth. “Kinematic and force analysis of articulated mechanical hands”. In: (1983).
- [44] Jingyi Xu et al. “Grasping posture estimation for a two-finger parallel gripper with soft material jaws using a curved contact area friction model.” In: *ICRA*. 2017, pp. 2253–2260.
- [45] Jingyi Xu et al. “6DFC: Efficiently Planning Soft Non-Planar Area Contact Grasps using 6D Friction Cones”. In: *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. 2020, pp. 7891–7897.
- [46] Jingyi Xu et al. “6DLS: Modeling Nonplanar Frictional Surface Contacts for Grasping Using 6-D Limit Surfaces”. In: *IEEE Trans. Robotics* (2021).
- [47] Andy Zeng et al. “Learning synergies between pushing and grasping with self-supervised deep reinforcement learning”. In: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. 2018, pp. 4238–4245.
- [48] Jialiang Zhao, Daniel Troniak, and Oliver Kroemer. “Towards Robotic Assembly by Predicting Robust, Precise and Task-oriented Grasps”. In: *Conf. on Robot Learning (CoRL)*. 2020.
- [49] Jiaji Zhou et al. “A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation”. In: *Int. Journal of Robotics Research (IJRR)* 37.2-3 (2018), pp. 249–265.