Simulation of Parallel-Jaw Grasping using Incremental Potential Contact Models

Chung Min Kim¹, Michael Danielczuk¹, Isabella Huang², Ken Goldberg¹

Abstract-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 simulator using Incremental Potential Contact (IPC) - a deformation model developed in 2020 for computer graphics - that models both the dynamics and the deformation of compliant jaw tips during grasping. IPC-GraspSim is evaluated using a set of 2,000 physical grasps across 16 adversarial objects where standard 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 backend), results suggest that IPC-GraspSim more accurately models real-world grasps, increasing F1 score by 9%. All data, code, videos, and supplementary material are available at https://sites.google.com/berkeley.edu/ipcgraspsim.

I. INTRODUCTION

Recent work in grasping objects with parallel-jaw grippers has revolved around learning to estimate high-quality grasps directly from RGB or depth images using a deep neural network [11, 18, 24, 31, 34, 37, 47]. These grasping networks are trained on large-scale datasets of labeled grasps generated either on physical systems [11, 18, 24, 37] or in simulation [17, 29, 31]. Simulation-based approaches can avoid time-intensive physical grasp attempts and have shown strong transfer to physical systems [31], but may suffer from noisy labels due to errors in modeling the interaction between gripper jaws and objects.

In practice, almost all parallel-jaw robot grippers use some form of compliant material on the jaw tips, from rubber coverings [35] to custom-made silicone rubber fingertips [31], as shown in Fig. 2. These compliant materials increase contact area by deforming around the grasped object to resist external tangential forces, increasing grasp robustness [40, 41, 44]. However, accurately modeling these compliant gripper jaws in simulation is difficult, as the interaction between the compliant material and the rigid object includes changing contact area, object pose, and frictional forces. Analytic modeling of compliant contacts, as a "soft finger" friction cone [5, 14] or as a discretized contact area mesh [8, 45], are limited to quasistatic analysis and thus do not incorporate dynamics. On the other hand, dynamic simulation of compliant jaws can be very sensitive to parameter choices: collision distances and timestep lengths must be manually tuned for different object



Fig. 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, predictions are incorrect; near-edge grasps are heavily dependent on gripper geometry, friction, and estimated object pose.

geometries and material parameters to avoid interpenetration or other physically unrealistic interactions between the jaws and object.

To address the need for modeling robust, soft contacts during dynamic parallel-jaw grasp motions, we present IPC-GraspSim, a grasping simulator using Incremental Potential Contact (IPC) [26]. IPC-GraspSim can accurately model grasp outcomes produced by a physical system, as shown in Fig. 1. IPC, developed in 2020 by the computer graphics community to simulate complex interactions of deformables, guarantees robust and intersection-free solutions for elastodynamics contact problems with friction, and supports variable material parameters, timestep lengths, and mesh resolutions.

¹The AUTOLab at University of California, Berkeley. {chungmin99, mdanielczuk, goldberg}@berkeley.edu ²Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA. isabella.huang@berkeley.edu



Fig. 2: Soft parallel-jaw grippers: (left to right) silicone rubber fingers in Dex-Net 4.0 [31], rubber tape in 6-DOF GraspNet [11], rubber fingertips in TossingBot [46]

This paper makes three contributions:

- IPC-GraspSim: an IPC-based simulator that models both friction and contact dynamics effects for compliant gripper jaws.
- 2) An Isaac Gym-based parallel-jaw grasping simulator with FleX backend that extends Eppner *et al.* [10] 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 85%, an improvement of 11% over analytic baselines and 9% over Isaac Gym.

II. RELATED WORK

A. 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.* [39], common metrics describe the ability of a grasp to resist external disturbances (force closure [3, 38]), be at equilibrium between contact forces and external forces [21], or allow the object to resist gravity (wrench resistance [30, 31]); 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, 27, 49] and can be extended to model non-planar contacts [44]. 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.

B. 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, 25, 28, 33], as it is important to model frictional contact dynamics simultaneously with gripper compliance: deformation affects surface area, which in turn affects object dynamics.

In particular, Isaac Gym is a GPU-enabled simulator favored by the robotics community for its speed. NVIDIA FleX engine integration has recently made deformables available for robotics environments, as evidenced by recent papers that model the SynTouch BioTac [36] 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, 29], (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].

C. 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 positionbased 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, Isaac Gym's softbody implementation, 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 [32]. 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 [26]. 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 [28, 41].

D. Incremental Potential Contact (IPC)

Incremental Potential Contact (IPC) [26] is a 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 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.

We leverage IPC as the basis of a grasping simulator that avoids the interpenetration issues described in Section II-C and more accurately models grasp outcomes observed on a physical system than existing quasistatic or dynamic simulator approaches.

III. 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.

A. 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.

B. 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



Fig. 3: Two grasp examples 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) grippers meet each other.

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 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.

C. 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), defined as $F1 = 2 \cdot \frac{AP \cdot AR}{AP + AR}$. F1 score is a common measure of binary classification performance for imbalanced datasets in computer vision [42].

IV. IPC-GRASPSIM

To simulate grasps, we adapt Incremental Potential Contact [26], as it guarantees intersection-free and inversion-free solutions for all parameter choices, as described in Section II-D. This model avoids the interpenetration artifacts faced by many simulator models [7, 28], 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, collision parameters need to be carefully tuned.

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.

A. 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 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 physical robots complete their gripper closing action when their maximum specified closing force is reached, contact force information is currently unavailable in the IPC simulator. To account for this, IPC-GraspSim approximates closing force on the target object by checking if A) the deformation energy described in Li et al. [26] 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 $\Psi_1 \approx \Psi_2$. (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 $\Psi_{th} = 5 \times 10^4$. Further details on Ψ_{th} , including sensitivity analysis, are in the supplement.

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.

V. IPC-GRASPSIM: 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 I (rectangular) and Table II (rounded), to study the effect of accurately modeling the gripper geometry.

A. 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.

Young's modulus	Friction coefficient μ										
E (Pa)	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 I: **Rectangular** jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 86%.

Young's modulus	Friction coefficient μ									
E (Pa)	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 II: **Rounded** jaw: results of parameter grid search over 3 training objects. $(E, \mu) = (10^8, 0.4)$ has the highest F1 score of 91%.

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

B. 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 in Figure 2. 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.

Compliant model	AP	AR	F1	Runtime (s)
Soft Point	0.76	0.55	0.64	0.01
REACH [8]	0.82	0.66	0.73	0.52
6DFC [45]	0.82	0.67	0.74	0.23
Isaac Gym	0.75	0.79	0.76	0.28
IPC-GraspSim	0.86	0.84	0.85	6.05

TABLE III: Average Precision (AP) and Average Recall (AR) for each model's grasp quality predictions for the 2000 grasps collected on the physical robot. Runtimes for analytic methods (Soft Point, REACH, 6DFC) are provided per-grasp, and runtimes for simulations (Isaac Gym, IPC-GraspSim) are provided per-frame (i.e., time to compute one timestep).

C. 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 $(10^7 - 10^8 \text{ Pa}, 0.46)$.

VI. EXPERIMENTS

A. 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.

B. 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 [28]. We reference the performance of the analytic models from experimental results in Xu *et al.* [45], and 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 V. Soft Point, REACH, and 6DFC are run on an Ubuntu 16.04 machine with a 12-core 3.7 GHz i7-8700k processor, Isaac Gym on an Ubuntu 20.04 machine with an Intel i7-6850K processor and NVIDIA Titan X GPU, and IPC-GraspSim on an Ubuntu 18.04 on a single core of a Dual 20-Core Intel Xeon E5-2698 v4 2.2 GHz.

C. Results

The results in Table III suggest that IPC-GraspSim outperforms all baselines across the dataset of 2,000 grasps and 16 objects, increasing F1 score by 11% over the best analytic model and 9% over Isaac Gym. We observe a 20% increase in average recall in IPC-GraspSim compared to grasp models that approximate compliance (Soft Point, REACH, 6DFC) and a 5% 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, 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.

In comparison to Isaac Gym, IPC-GraspSim increases average precision by 11%, average recall by 5%, and F1 score by 9%. The performance of IPC-GraspSim highlights the importance of intersection-free and robust modeling of soft contacts for compliant gripper modeling. First, poorlyresolved contacts may cause grasped objects to move unpredictably (e.g., float or swing into air, as shown in Figure 4C-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.

D. Failure Mode Analysis

Table IV 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 0.83 as compared to 0.76 for the best-performing among the analytic methods and 0.68 for Isaac Gym.

We observe a false positive prediction for grasp 7, where the 3D printed plastic part of the gripper jaws bend or 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 and do not account this bending phenomenon.

We also observe a false negative prediction for grasp 2, where the right jaw contacts the vase first and rotates the vase out of the gripper. This phenomenon is also observed in Figure 4A, where the simulated grasp fails as the object is

Grasp sets (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	0.0	1.0	0.0	0.0	1.0	0.4	0.8	1.0	1.0	1.0	1.0	1.0	1.0	0.6	0.0	0.0	0.0	0.4	0.0	0.0	1.0	1.0	0.0	0.0
Physical Experiments	0.0	1.0	1.0	0.0	0.0	1.0	0.8	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	0.2	0.0	0.0	0.4	1.0	0.8	0.0	0.0

TABLE IV: Grasp robustness values predicted by IPC-GraspSim (top row) and observed from physical experiments (bottom row) using an ABB YuMi with a compliant parallel-jaw gripper on the "vase" object shown in Fig. 1. 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, which suggests that IPC-GraspSim can reliably predict physical grasp outcomes. Example true positive (grasp 2), true negative (grasp 0), false positive (grasp 7), and false negative (grasp 1) grasps are shown in Fig. 1.



Fig. 4: Four grasps (columns) with three simulators: (Top to bottom) IPC-GraspSim, Isaac Gym, 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 IPC-GraspSim, rigid jaws have sufficient contact area to grasp "pawn". In Isaac Gym, incorrect contact forces push the object up and away from the plane. Physical grasp fails as jaw cantilevering decreases contact area. (D) False negative: in IPC-GraspSim, left gripper jaw pushes object out of the gripper. For Isaac Gym, the object is pushed upwards, with a similar issue as described in (C). Videos of all simulator grasp attempts can be found on the website.

not contacted by the right gripper jaw. 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.

VII. DISCUSSION

In this paper, we present IPC-GraspSim, an intersectionfree, robust grasp simulator, and evaluate its performance on a physical grasp dataset generated with compliant parallel-jaw grippers. 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. In future work, we will accelerate the simulator by using Medial-IPC [23], which has demonstrated speed performance improvements up to 110x times faster than the original IPC implementation and explore the ability of IPC-GraspSim to generalize to different jaw materials and geometries. We will also explore the effect of training grasp prediction neural networks using the predicted grasp outcomes from IPC-GraspSim.

ACKNOWLEDGEMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, and the CITRIS "People and Robots" (CPAR) Initiative. The authors were supported in part by donations from Google, Siemens, Toyota Research Institute, Autodesk, Honda, Intel, and Hewlett-Packard. Michael Danielczuk is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1752814. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors. We thank our colleagues who provided helpful feedback and suggestions, in particular Jeffrey Ichnowski who introduced Incremental Potential Contact (IPC) work to us and AUTOLab. We also thank Minchen Li and Zachary Ferguson for support in using IPC.

REFERENCES

- [1] A. Allshire, M. Mittal, V. Lodaya, V. Makoviychuk, D. Makoviichuk, F. Widmaier, M. Wüthrich, S. Bauer, A. Handa, and A. Garg, "Transferring dexterous manipulation from gpu simulation to a remote real-world trifinger," in *Conf. on Robot Learning (CoRL)*, 2021.
- [2] J. Bender, M. Müller, M. A. Otaduy, M. Teschner, and M. Macklin, "A survey on position-based simulation methods in computer graphics," in *Computer graphics forum*, Wiley Online Library, vol. 33, 2014, pp. 228–251.
- [3] A. Bicchi, "On the closure properties of robotic grasping," Int. Journal of Robotics Research (IJRR), vol. 14, no. 4, pp. 319–334, 1995.
- [4] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, vol. 1, 2000, pp. 348–353.
- [5] M. Ciocarlie, C. Lackner, and P. 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] M. Ciocarlie, A. Miller, and P. Allen, "Grasp analysis using deformable fingers," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005, pp. 4122–4128.
- [7] E. Coumans and Y. Bai, Pybullet, a python module for physics simulation for games, robotics and machine learning, http:// pybullet.org, 2016–2021.
- [8] M. Danielczuk, J. Xu, J. Mahler, M. Matl, N. Chentanez, and K. Goldberg, "Reach: Reducing false negatives in robot grasp planning with a robust efficient area contact hypothesis model," in *Int. S. Robotics Research (ISRR)*, 2019.
- [9] C. Eppner, A. Mousavian, and D. 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] —, "Acronym: A large-scale grasp dataset based on simulation," in Proc. IEEE Int. Conf. Robotics and Automation (ICRA), 2020.
- [11] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-Ibillion: A large-scale benchmark for general object grasping," in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11444–11453.
- [12] S. Goyal, A. Ruina, and J. Papadopoulos, "Planar sliding with dry friction part 1. limit surface and moment function," *Wear*, vol. 143, no. 2, pp. 307–330, 1991.
- [13] M. Guo, D. V. Gealy, J. Liang, J. Mahler, A. Goncalves, S. McKinley, J. A. Ojea, and K. Goldberg, "Design of parallel-jaw gripper tip surfaces for robust grasping," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 2831–2838.
- [14] H. Hertz, Miscellaneous papers. Macmillan, 1896.
- [15] R. D. Howe, I. Kao, and M. 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] I. Huang, Y. Narang, C. Eppner, B. Sundaralingam, M. Macklin, T. Hermans, and D. Fox, "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] E. Johns, S. Leutenegger, and A. 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] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, *et al.*, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *Conf. on Robot Learning (CoRL)*, 2018, pp. 651–673.
- [19] I. Kao and M. R. Cutkosky, "Quasistatic manipulation with compliance and sliding," *Int. Journal of Robotics Research (IJRR)*, vol. 11, no. 1, pp. 20–40, 1992.
- [20] I. Kao and F. Yang, "Stiffness and contact mechanics for soft fingers in grasping and manipulation," *IEEE Trans. Robotics and Automation*, vol. 20, no. 1, pp. 132–135, 2004.
- [21] J. Kerr and B. Roth, "Analysis of multifingered hands," Int. Journal of Robotics Research (IJRR), vol. 4, no. 4, pp. 3–17, 1986.
- [22] G. Kloosterman, "Contact methods in finite element simulations," Undefined, Ph.D. dissertation, Dec. 2002.
- [23] L. Lan, Y. Yang, D. Kaufman, J. Yao, M. Li, and C. Jiang, "Medial ipc: Accelerated incremental potential contact with medial elastics,"

ACM Transactions on Graphics (TOG), vol. 40, no. 4, pp. 1–16, 2021.

- [24] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *Int. Journal of Robotics Research (IJRR)*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [25] B. León, S. Ulbrich, R. Diankov, G. Puche, M. Przybylski, A. Morales, T. Asfour, S. Moisio, J. Bohg, J. Kuffner, *et al.*, "Open-grasp: A toolkit for robot grasping simulation," in *Proc. IEEE Int. Conf. on Simulation, Modeling, and Programming of Autonomous Robots (SIMPAR)*, Springer, 2010, pp. 109–120.
- [26] M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman, "Incremental potential contact: Intersection-and inversion-free, large-deformation dynamics," ACM transactions on graphics, 2020.
- [27] Y. Li and I. 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.
- [28] J. Liang, V. Makoviychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "Gpu-accelerated robotic simulation for distributed reinforcement learning," in *Conf. on Robot Learning (CoRL)*, 2018, pp. 270–282.
- [29] J. Mahler and K. Goldberg, "Learning deep policies for robot bin picking by simulating robust grasping sequences," in *Conf. on Robot Learning (CoRL)*, 2017, pp. 515–524.
- [30] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "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.
- [31] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, eaau4984, 2019.
- [32] G. Marckmann and E. Verron, "Comparison of hyperelastic models for rubber-like materials," *Rubber chemistry and technology*, vol. 79, no. 5, pp. 835–858, 2006.
- [33] A. T. Miller and P. K. Allen, "Graspit! a versatile simulator for robotic grasping," *IEEE Robotics & Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [34] D. Morrison, P. Corke, and J. Leitner, "Learning robust, real-time, reactive robotic grasping," *Int. Journal of Robotics Research (IJRR)*, vol. 39, no. 2-3, pp. 183–201, 2020.
- [35] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *Proc. IEEE Int. Conf.* on Computer Vision (ICCV), 2019, pp. 2901–2910.
- [36] Y. Narang, B. Sundaralingam, M. Macklin, A. Mousavian, and D. Fox, "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.
- [37] L. Pinto and A. 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.
- [38] E. Rimon and J. 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.
- [39] M. A. Roa and R. Suárez, "Grasp quality measures: Review and performance," Autonomous robots, vol. 38, no. 1, pp. 65–88, 2015.
- [40] J.-P. Roberge, W. Ruotolo, V. Duchaine, and M. Cutkosky, "Improving industrial grippers with adhesion-controlled friction," *IEEE Robotics & Automation Letters*, vol. 3, no. 2, pp. 1041–1048, 2018.
- [41] A. Rocchi, B. Ames, Z. Li, and K. Hauser, "Stable simulation of underactuated compliant hands," in *Proc. IEEE Int. Conf. Robotics* and Automation (ICRA), 2016, pp. 4938–4944.
- [42] T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets," *PloS one*, vol. 10, no. 3, e0118432, 2015.
- [43] J. K. Salisbury and B Roth, "Kinematic and force analysis of articulated mechanical hands," 1983.
- [44] J. Xu, T. Aykut, D. Ma, and E. Steinbach, "6dls: Modeling nonplanar frictional surface contacts for grasping using 6-d limit surfaces," *IEEE Trans. Robotics*, 2021.
- [45] J. Xu, M. Danielczuk, E. Steinbach, and K. Goldberg, "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] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossingbot: Learning to throw arbitrary objects with residual physics," *IEEE Trans. Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [47] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "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] J. Zhao, D. Troniak, and O. Kroemer, "Towards robotic assembly
- [48] J. Zhao, D. Troniak, and O. Kroemer, "Towards robotic assembly by predicting robust, precise and task-oriented grasps," in *Conf. on Robot Learning (CoRL)*, 2020.
 [49] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell, "A convex
- [49] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell, "A convex polynomial model for planar sliding mechanics: Theory, application, and experimental validation," *Int. Journal of Robotics Research* (*IJRR*), vol. 37, no. 2-3, pp. 249–265, 2018.