How Firm Should a Grasp Be?

Matthew Beveridge
Columbia University
beveridge@cs.columbia.edu

Shree K. Nayar
Columbia University
nayar@cs.columbia.edu

Abstract: An ideal robot grasp is firm enough to securely handle an object, yet gentle enough to avoid damaging it. Achieving this balance requires knowledge of the object's material properties, such as its density, elasticity, and surface friction. These properties, however, are seldom precisely known *a priori*. In this work, we propose a visuotactile approach to estimating material properties in real time, during the process of grasping. Our method uses these estimated properties to determine the minimum grasp force required to handle the object. We contribute a new dataset of real-world objects (fruits and vegetables) with measured physical properties (shape, mass, elasticity, and friction), which we use to construct our force estimation model via simulations. We experimentally validate our approach to grasp force control using a robot with a parallel-jaw gripper. We demonstrate our system's ability to gently grasp a wide variety of objects, in each case adapting to their unique physical properties.

Keywords: Material Recognition, Deformable Object Grasping, Tactile Sensing, Force Control

1 Gentle Grasping Requires Material Perception

Given an object and a desired interaction with it, the ideal grasp should exert no more force on the object than necessary to complete the task. While the notion of damage is often considered in catastrophic terms (*e.g.*, a glass shattering), it can also be more subtle. For example, with excessive force, a fruit may bruise, a cardboard box might dent, or a plastic toy could crack or get scratched. In addition, there is a temporal element to damage – repeated interactions without care for an object's fragility can cause cumulative wear over time. In this work, we are interested in the *gentle grasping* of objects, where the goal is to minimize the potential for damage during a grasp.

Since the material properties of an object are seldom precisely known *a priori*, how do we, humans, determine how firm its grasp should be? Dynamic grasp adaptation is intuitive to us [1, 2] – we use our vision to obtain priors on the object's material composition [3–5], and rely on tactile feedback to refine these priors while grasping it [6, 7]. Since some of the material cues we need to understand an object's physics cannot be known from vision alone, we estimate them via touch during the nascent stage of a grasp. We use these estimates to continuously adjust the forces we apply to the object and are thus able to gracefully handle it without conscious effort.

For a robot to gently grasp an object, it must either know the object's properties beforehand or be able to infer them during the grasp. Unless the object is precisely engineered or measured in detail, the best the robot has to work with is a rough approximation of its properties. One method to refine this approximation is for the robot to interact with the object before attempting a grasp. That is, the robot "probes" the object with a sequence of exploratory actions to estimate, for instance, its hardness [8, 9]. Such an approach is time-consuming and hence not suitable for applications where speed is critical. Instead, we propose to leverage the dynamic, ephemeral tactile signals that are sensed *during a grasp* to estimate the object's properties in real time. The robot is therefore able to adjust the force it applies to the object on-the-fly, *in situ*.

In a typical robotic task, there are four initial phases, as illustrated in Fig. 1. The first is *observe*, where, from a distance, the robot can visually identify the class of the object, its shape, and its location. During the *approach* phase, the robot positions itself at a suitable grasp location, determined from the object's geometry and pose. Next, we have *grip*, where the gripper squeezes the object with a specified force to take hold of it. Finally, the robot *lifts* the object and executes a goal task, such as moving it to a new location or mating it with another object. In our work, we focus on the *observe*, *grip*, and *lift* phases and assume that for *approach* a reasonable grasp pose is provided by an existing planner [10, 11]. The *observe* phase reveals the class of the object ("pepper" in Fig. 1),

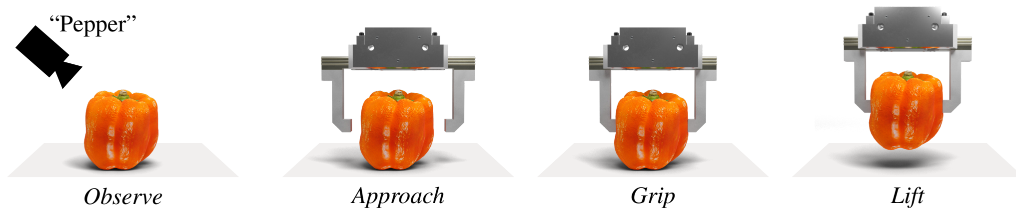

Figure 1. **Phases of a Grasp.** We consider a grasp to comprise four stages: observe, approach, grip, and lift. Our work exploits the short-lived (fraction of a second) haptic signals during the grip and lift phases to refine class-level material properties (from observation) to obtain instance-level estimates. During this refinement, we continuously adjust the applied force to ensure it is just above the minimum force needed to handle the object.

giving us a prior on its properties (mass, elasticity, friction, *etc*.). During the *grip* and *lift* phases, we use haptic feedback to continuously refine these priors and control the force applied to the object.

While haptic signals encode an object's material properties, different combinations of material parameters could lead to the same haptic signals. Therefore, material estimation can be viewed as a challenging inverse problem. To develop our solution to this inverse problem, we have chosen the domain of fruits and vegetables. This is an interesting domain as two samples, even when of the same class, can be expected to differ in their properties (*e.g*., a raw and a ripe avocado). We first collected a dataset, called *Squash*, that includes material measurements and 3D scans of 270 samples across 52 different classes of fruits and vegetables. [1] Using *Squash*, we derived class-level property distributions, which reduce the search space for instance-level properties.

Using *Squash*, we simulated numerous object grasps, with augmentations of material properties to span real-world variability. We trained a sequence model on the evolution of the simulated haptic data (*i.e*., tactile images and load cell measurements) to distill instance-level material properties from class distributions. We have integrated our visuotactile model for force control into a robot's operating system, where it runs at 250 Hz. We used a 6 DoF arm and a parallel-jaw gripper with GelSight tactile sensors to validate our approach. Our system achieved more than a 2x reduction in excess applied force for a successful grasp compared to non-adaptive baselines.

2 Related Work

Class-Level Material Recognition from a Distance. For grasping, it is beneficial to recognize an object's material class (*e.g*., peach) passively during the *observation* stage. However, precise class-level visual recognition (*i.e*., peaches vs. nectarines) requires elaborate optical measurements, such as the reflectance distribution function (BRDF) [12–14], which is hard in dynamic, cluttered robotics settings. Hence, class recognition is most often done from appearance alone [15–19], using filter banks, statistical analysis, and deep learning. These approaches do not directly relate appearance to material properties, which has been the goal of more recent work [20–22]. In our work, we rely on appearance to identify class-level material properties during *observation*, while recognizing that additional sensing modalities are needed to resolve instance-level variation within a class.

Instance-Level Material Recognition through Excitation. During the *grip* and *lift* stages, haptic feedback provides material-related signals for the specific object of interest (*e.g*., *this* peach has *this* hardness). In the general context of measuring physical properties, exciting objects with ultrasound and physical impact have been shown to reveal useful material cues [23–25]. In the realm of robotics, vision-based tactile sensors have been used to provide detailed insights related to contact mechanics. For example, GelSight [26–28] and DIGIT [29] have been used for hardness estimation [30] and shear/slip detection [31]. Beyond optical sensors, piezoresistive [32, 33] and magnetic tactile skins [34], which measure pressure distributions, have been integrated into grippers for fine-grained visuotactile manipulation [35, 36]. Cross-modal methods have also been developed to relate visual appearance to tactile properties [21, 37–40]. In our work, we use GelSight sensors to image contact geometry, shear, and slip during a grasp. Our framework for material-aware grasping, however, is applicable across a variety of tactile sensing modalities.

[1] The dataset is available on the *Squash* webpage.

Approaches to Gentle Grasping. A classical approach to avoiding object damage is to regulate contact force via impedance control [41], or to adjust grasp force reactively using tactile feedback [42–44], *e.g.*, upon detecting incipient slip [45, 46]. For this purpose, tactile sensors have also been integrated into soft robotic hands due to their passive compliance [47, 48]. Purely vision-based methods [49–55] have also demonstrated successful deformation-aware grasping, but are shown to be more effective when fused with tactile sensing [56, 57]. More recently, human and robot demonstrations have enabled end-to-end gentle grasping of objects [8, 9, 36, 58, 59], and diffusion-based policies conditioned on tactile signals have been used for force-aware manipulation of fragile objects [60]. Nevertheless, these methods are not real-time, do not expose inferences about the object's material, and only yield predictions of an "acceptable" grasp force based on heuristics.

3D Object Datasets. Several large-scale 3D object datasets have been developed for robotic manipulation [61–65] and broader 3D vision tasks [66–68]. These datasets provide geometry and appearance but do not include measured material properties. The recent work by Cao and Kalogerakis [69] uses annotated simulated objects with material parameters, where the parameter values are assigned using material-category labels rather than measurements made on individual instances. For deformable objects, smaller datasets exist that include detailed annotations of the deformation process through finite element analysis [49, 70]. One of our contributions is *Squash*, a large dataset that pairs 3D scans of fruits and vegetables with per-instance measurements of material properties (elasticity, mass, and surface friction). It is the first 3D object dataset specifically designed for deformable object simulation with measured properties. Beyond its utility in robotics, we expect *Squash* to be a valuable resource for simulation and rendering in computer graphics and animation.

3 The Physics of Grasping a Material

If we know the material properties of an object, then what minimal grasp force should we use to lift it? Naïvely, the force F required to suspend a rigid body of mass m at equilibrium is $F = mg$, *i.e.*, Newton's second law. If we assume that the force due to friction is independent of the mass and the contact area, we would get the familiar modification: $\mu F = mg$, where μ is the coefficient of friction. Coulomb observed that the effect of friction is primarily due to the interlocking of surface asperities, which he assumed to be infinitesimally small. This, however, is an approximation. Bowden and Tabor [71] noted that friction capacity can vary under pressure, thereby acting as an additional load, especially when the interacting bodies are soft. The effect of this additional load is visualized in Fig. 2, where we see that soft objects require a lower grasp force than rigid ones, since they deform and create a larger contact area under the same load. Formally, this yields a new equilibrium:

$$\mu F + \tau A_{\text{real}} = mg, \qquad (1)$$

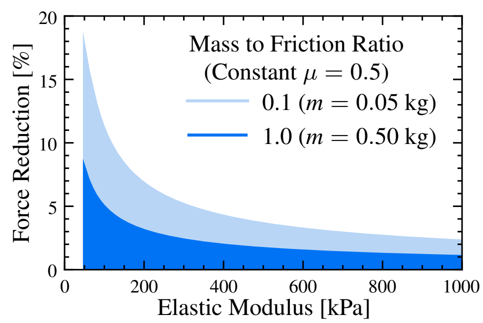

Figure 2. **Soft Objects Require Less Grasp Force.** For a soft object, the contact area A grows more rapidly with applied force F than it does for a hard object. As A grows, an additional load, τA, where τ is the object's material shear strength, supplements F to resist the effect of gravity (Eq. (1)). The theoretical reduction in minimum-required grasp force due to this additional load is shown here as a function of the elastic modulus. The object and its shear strength τ are kept constant in the plot, but the ratio of mass m to friction μ is varied. The greatest benefit (*i.e.*, the possible reduction in grasp force) is for light, pliable objects.

where τ is the interfacial shear strength, and A_{real} is the contact area over which surface asperities interact. For rigid bodies, A_{real} is considerably smaller than the apparent contact region A_{app} (on the macroscale), as only the peaks of the asperities interlock [72].

If the object or the gripper is soft, we can say $A_{\text{real}} \approx A_{\text{app}}$, since the elastic deformation fills the cavities between the summits of surface asperities. In our context, A_{app} can be directly observed through tactile sensing. Even so, F and A_{app} are directly coupled, and without a mapping between them, we must *search* for the force that results in the equilibrium given by Eq. (1). To *predict* this force of equilibrium, we must also predict how the apparent contact area will change with respect to the load force. Finite element analysis (FEA) can accurately estimate this contact area, assuming

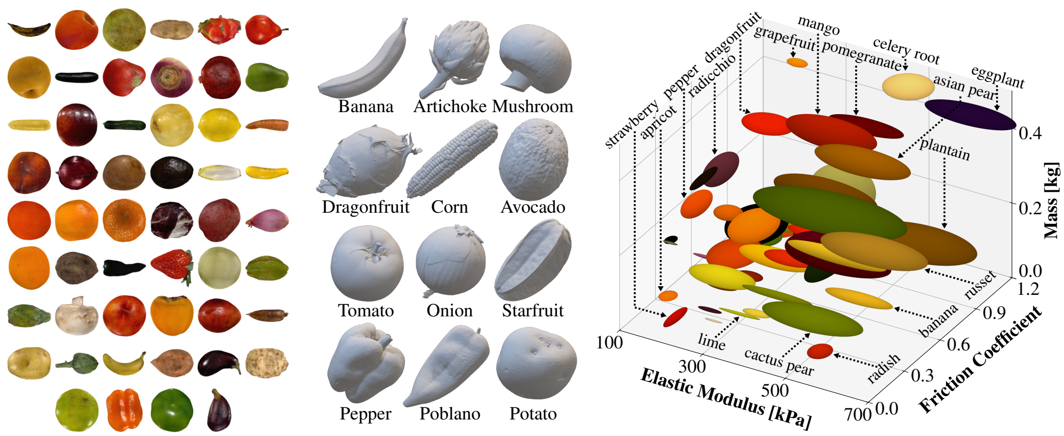

(a) 52 Material Classes (b) Example 3D Scans (c) Mechanical Property Distributions

Figure 3. **The *Squash* Dataset.** A dataset of real-world 3D object meshes and material property measurements. (a) *Squash* is composed of 270 instances across 52 classes of produce (fruits and vegetables). (b) We scanned each object using structured light to measure its geometry and texture at high resolution. (c) We additionally measured the material properties of each object (elastic modulus, mass, volume, and surface friction), where each ellipsoid represents the distribution of the properties within a class. *Squash* is a first-of-its-kind dataset that captures the natural variability of object properties, which can be used for physically-accurate simulation.

the material properties of the object and gripper are known, *a priori*. In our case, property estimates change during the course of a grasp, necessitating an expensive FEA reevaluation on each update, which would limit the rate at which we can adapt grasp force.

Instead, by using a Hertzian contact approximation [73, 74], A_{app} can be modeled using a closed-form expression. If we approximate the fingertips of a robotic gripper as near-planar, and an object as locally spherical, then applying a normal force F yields:

$$A_{\text{app}}(F) = \pi \left(\frac{3FR^*}{4E^*}\right)^{2/3}, \text{ given } \frac{1}{E^*} = \frac{1-\nu_{\text{obj}}^2}{E_{\text{obj}}} + \frac{1-\nu_{\text{grip}}^2}{E_{\text{grip}}} \text{ and } R^* = \frac{R_{\text{obj}}\,R_{\text{grip}}}{R_{\text{obj}}+R_{\text{grip}}}, \tag{2}$$

where E^* is the effective elastic modulus between the object and the gripper, R^* is the effective radius of curvature between the object and the gripper at the contact interface, and ν is the Poisson ratio. Given the mass m, elasticity E, surface friction μ, and interfacial shear strength τ of an object, we can find A_{app} using Eq. (2) and substitute in Eq. (1) to estimate the minimum force needed to lift the object. But, for an arbitrary object, we at best know a range of plausible values for (E, m, μ, τ) based on its class. Our goal, therefore, is to develop a learning-based approach to refine (E, m, μ, τ) through temporal changes in haptic feedback, which we then use to servo-control the applied force.

4 *Squash*: A Dataset of 3D Object Assets and their Physical Properties

To learn our visuotactile model for gentle grasping, we require a dataset of deformable objects with known physical properties. Deformable objects are interesting in two respects. First, they are more prone to damage due to a grasp that is too firm (*e.g.*, a peach will bruise). Second, as outlined in Sec. 3, the contact area between the gripper and the object changes as they both deform, thereby requiring continuous control of the minimum grasp force. We have developed *Squash* (Fig. 3), a comprehensive dataset of 3D scans and physical property measurements of various fruits and vegetables. *Squash* contains 270 instances of produce spread across 52 classes (capturing intraclass variability). All 52 classes are shown in Fig. 3a, and a few examples of the 3D scans are shown in Fig. 3b. Using the methods described below, we measured the elastic modulus, mass, volume, and surface friction of each instance. A visualization of the measured properties is shown in Fig. 3c, where each ellipsoid represents the distribution of the properties within a class.

Characterizing the Properties of an Object. For each object, we first measured its mass and Shore hardness using a durometer [75, 76]. From its 3D scan, we then computed its volume. Next, using a robot arm and a parallel-jaw gripper affixed with load cells and tactile fingertips (detailed in the

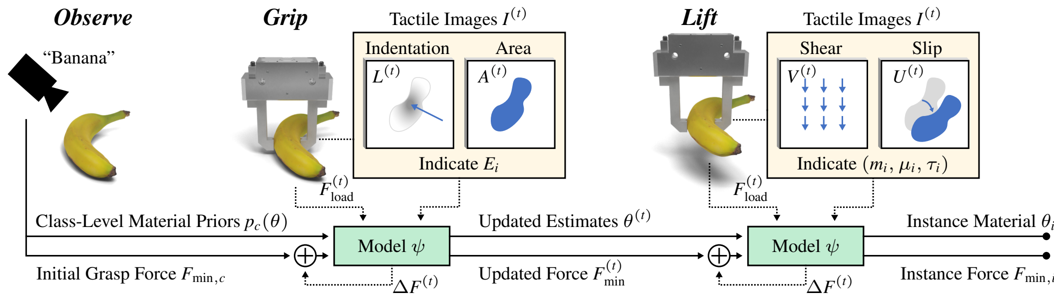

Figure 4. **Framework for Gentle Grasping.** From *observation*, we identify the object's class to get a class-conditioned material prior and an initial minimum-required grasp force estimate. Upon first contact, our visuo-tactile likelihood model ψ (detailed in Fig. 5) begins to refine these priors into instance-level posteriors from temporal tactile feedback. The minimum-required grasp force is continuously reestimated in the process, and applied in real time. The result is a gentle grasp that adapts to the object instance, *in situ*.

supplemental material), we measured the object's elasticity and friction capacity by performing a compression test and a tribological test, respectively. In the compression test, with the object at rest, we recorded the applied force F_{load}, contact area A, and indentation into the object L, while closing the gripper jaws on the object. We modeled the local surface of the object as quasi-spherical and the gripper jaws as parallel plates. Then, for any deformation of the object, the effective modulus E^* follows from Eq. (2): $E^* = \frac{3F}{4L}\sqrt{\frac{\pi}{A}}$, since from Hertzian mechanics, the effective curvature of the contact interface is $R^* = \frac{A}{\pi L}$. We have many estimates of effective modulus E^* (during closure), and we average them for our final estimate of E^*. We then recover the object's elastic modulus E from E^* by separately measuring the elastic modulus of the tactile sensor. We report E as a function of the Poisson ratio ν, since we cannot directly observe the equatorial deformation of the object.

For the tribological test, we begin with the object lifted, then increase the gripper opening until slip is detected by the tactile sensor. Each trial yields one $(F_{\text{slip}}, A_{\text{slip}})$ pair, where F_{slip} is the applied force just before slip, and A_{slip} is the contact area at that onset. These measurements are related through Eq. (1): $\mu F_{\text{slip}} + \tau A_{\text{slip}} = mg$. We collected several trials per object, each at a different location on its surface, and solved the over-determined linear system in (μ, τ) by least squares. By sampling distinct locations on the object, we capture local deviations in surface curvature. This gives us different couplings between F_{slip} and A_{slip}, enabling the decoupling of μ and τ during estimation. Since *Squash* includes a large sampling of shapes, masses, elasticity parameters, and friction coefficients, we can use it to generate diverse grasp simulations for model training.

5 Material-Aware Robotic Grasping

Fig. 4 illustrates our closed-loop grasping framework for estimating material properties $\theta = (E, m, \mu, \tau)$ to determine the minimum necessary grasp force F_{min} for an object. The material prior distribution $p_c(\theta) \equiv p(\theta \mid c)$ for the object class c, identified during *observation*, yields an initial estimate of the minimum grasp force $F_{\text{min},c}$ via Eq. (1). The *visuotactile likelihood model* ψ then refines this class prior during contact with the object, yielding an online posterior belief over material properties at each time t, denoted as $\theta^{(t)}$. As the robot begins to grip the object, tactile images $I^{(t)}$ are used to update the elasticity estimate $E^{(t)}$ from temporal changes in indentation $L^{(t)}$ and contact area $A^{(t)}$ relative to the applied load $F_{\text{load}}^{(t)}$.[2] As the robot begins to lift the object, shear $V^{(t)}$ and slip $U^{(t)}$ in the tactile images are used to update the estimated mass $m^{(t)}$ and friction parameters $\mu^{(t)}$ and $\tau^{(t)}$. The model ψ therefore continuously refines the material estimate $\theta^{(t)} = (E^{(t)}, m^{(t)}, \mu^{(t)}, \tau^{(t)})$, and the controller accordingly adjusts the applied grasp force $F_{\text{load}}^{(t)}$ to achieve the current minimum-force estimate $F_{\text{min}}^{(t)}$. As additional tactile observations arrive, the posterior approaches the instance-level material properties θ_i, yielding the desired object-specific force estimate $F_{\text{min},i}$ for gentle grasping.

Learning Material from Simulated Grasps. To develop our visuotactile likelihood model ψ, we first designed a simulator to generate high-fidelity tactile data for training. We built upon IPC [77–

[2] While our implementation uses GelSight tactile images, the model ψ can utilize other tactile sensor outputs, *e.g.*, binary contact maps and pressure maps. For our purposes, we refer to all these outputs as tactile images.

79] for collision physics, which has been shown to accurately model real-world grasps of deformable objects [80, 81]. We synthesized 200k training grasp sequences using a parallel-jaw gripper and the object meshes in *Squash*. Object properties were sampled from the distributions in *Squash* and supplemented by Latin Hypercube sampling. Each sequence contains temporal measurements of each jaw's tactile image $I^{(t)}$ and applied normal load $F^{(t)}_{\text{load}}$. We then train ψ, shaded green in Fig. 5, to encode these temporal haptic observations. At each timestep t, the tactile image $I^{(t)}$ is embedded by an image encoder, paired with $F^{(t)}_{\text{load}}$, and passed to a closed-form continuous-time neural network (CfC) [82] to produce a hidden state $\boldsymbol{h}^{(t)}$. A CfC is chosen for its ability to model complex dynamics using ordinary differential equations, its ability to handle non-uniformly sampled robot data, and its inference speed. We trained ψ with a hybrid contrastive-regression loss over the simulated grasps. At inference, $\boldsymbol{h}^{(t)}$ parameterizes the likelihood used to score candidate particles $\theta^{(t)}_j$ in a particle filter, producing an online posterior belief over material properties. Further details are in the supplemental material.

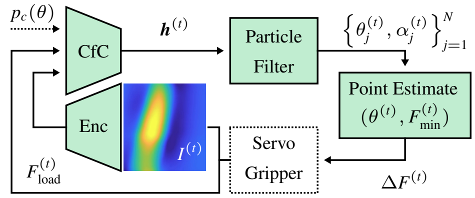

Figure 5. **Visuotactile Model ψ (shaded green) for Material Estimation and Force Control.** We start with the class prior $p_c(\theta)$ for an object's material θ. Tactile images $I^{(t)}$ and the applied load $F^{(t)}_{\text{load}}$ are then encoded by the continuous-time sequence model (CfC) into the hidden state $\boldsymbol{h}^{(t)}$. The particle filter uses this representation to update the weighted material belief $\{\theta^{(t)}_j, \alpha^{(t)}_j\}^N_{j=1}$. The resulting point estimate $\theta^{(t)}$ determines the current estimate of minimum-required grasp force $F^{(t)}_{\text{min}}$, which is used as the target force by the gripper servo-controller. The entire control loop (including ψ) runs at 250 Hz.

Gentle Grasping from Temporal Observations. Using ψ and sequential observations $(I^{(0:t)}, F^{(0:t)}_{\text{load}})$, we update the posterior belief over material properties θ as:

$$p(\theta \mid I^{(0:t)}, F^{(0:t)}_{\text{load}}) \propto p_\psi(I^{(t)}, F^{(t)}_{\text{load}} \mid \theta)\, p(\theta \mid I^{(0:t-1)}, F^{(0:t-1)}_{\text{load}}), \tag{3}$$

where $p_\psi(\cdot \mid \theta)$ is the learned likelihood induced by ψ. We represent this belief with particles,

$$p(\theta \mid I^{(0:t)}, F^{(0:t)}_{\text{load}}) \approx \sum\nolimits_j \alpha^{(t)}_j \delta(\theta - \theta^{(t)}_j), \quad \alpha^{(t)}_j \propto \alpha^{(t-1)}_j p_\psi(I^{(t)}, F^{(t)}_{\text{load}} \mid \theta^{(t)}_j), \tag{4}$$

where $\alpha^{(t)}_j$ is the weight of particle $\theta^{(t)}_j$. Since particles are resampled during filtering, both particle locations and weights are indexed by time. Each particle represents a candidate contact model and grasp feasibility condition [83, 84]. While lifting the object, feasibility reduces to whether the frictional support exceeds the object's weight. The minimum required grasp force is therefore:

$$f(\theta) = \min\nolimits_{F \geq 0} \{F : \mu F + \tau A(F; E) \geq mg\}, \tag{5}$$

following from Eq. (1), where $A(F;E)$ is given by Eq. (2). We compute a force hypothesis $F^{(t)}_{\text{min},j} = f(\theta^{(t)}_j)$ for each particle, and the weighted set $\{F^{(t)}_{\text{min},j}, \alpha^{(t)}_j\}$ represents the current uncertainty in the required grasp force. As additional haptic observations arrive, posterior uncertainty over θ contracts, producing a tighter estimate of the minimum force needed to lift the object. Thus, ψ refines class-level material priors into instance-level estimates, enabling gentle, material-aware grasping, *in situ*.

6 Experiments

Evaluation via Simulations. We simulated 5K objects using the simulator described in Sec. 5 and the meshes in *Squash*. We procedurally deformed the meshes to create realistic deformations, and sampled material properties from the distributions in *Squash*. We then compared our approach to two baseline grasp force policies: (1) *Naïve*, where the commanded grasp force is always 3 N (sufficient to lift the largest object in *Squash*) regardless of the object's class, and (2) *Class Prior*, where grasp force is determined using Eq. (1) from the class's average material properties in *Squash*. The results are summarized in Tab. 1, where we report the performance of force control and material property estimation. Our material-adaptive approach significantly outperforms both baselines in terms of grasp force accuracy through visuotactile estimation of the object's material properties. Notably, the *Naïve* approach results in a large excess force for many objects, while the *Class Prior* approach can lead to both over- and under-exertion due to the variability of material properties within a class. In contrast, our method achieves a much lower excess force on average, demonstrating its ability to adapt to the specific material properties of each object and maintain a high success rate.

Table 1. **Efficacy of Material-Aware Robot Grasping.** We compare our material-aware approach for gripper force control with two standard approaches using 5k simulated produce items with high variability. "Naïve" uses a force of 3 N in each gripper jaw, which just exceeds the threshold to successfully lift the largest object in *Squash*. "Class Prior" uses the class's average material properties to determine the grasp force using Eq. (1). In contrast, our approach refines instance-level mechanical properties to yield improved fine-grained force control. The best results are in bold, and second best are underlined. Ablations are reported in the supplemental material.

	Mean Absolute Percentage Error (MAPE) [%] ↓				Excess Force ↓			Success ↑
Approach	F	E	m	(μ, τ)	Overload [N]	Work [mJ]	Impulse [N s]	Rate [%]
Naïve (3 N)	249.5±128.7	—	—	—	2.01±0.41	3.32±2.04	2.15±1.28	**99.6±6.5**
Class Prior	207.0±221.6	107.4±25.4	150.3±213.2	98.7±23.2	1.63±1.42	3.63±5.76	1.95±2.58	89.0±31.3
Ours	**38.8±37.2**	**56.3±22.6**	**49.3±24.1**	**49.8±21.5**	**0.39±0.26**	**2.25±5.37**	**0.66±0.90**	99.1±9.2

Grasping Real Fruits and Vegetables. We next assembled a test set of 29 real fruits and vegetables. Their material properties were measured (using the methods in Sec. 4) as a ground truth reference and used to determine a theoretically optimal minimum grasp force. We then performed grasping trials with our material-aware approach using the robot hardware shown in Fig. 6a, and compared the applied grasp force to the theoretical minimum. In Figs. 6b and 6c, we show the evolution of the applied grasp force and estimated material properties over the duration of grasping a nectarine. As the robot progresses through the *grip* phase, the estimate of elasticity is refined. During the *lift* phase, mass and friction parameters are updated. The force and material estimates converged to stable values quickly over the course of the grasp, enabling the robot to apply a force that never significantly exceeds the theoretically optimal force for the nectarine. Fig. 6d compares the grasp forces of our method with those of the *Naïve* and *Class Prior* baselines, for a wide variety of fruits and vegetables. As seen in Fig. 6e, the distribution of excess forces for our approach is more tightly concentrated around zero, compared to the baselines. Finally, in Figs. 6f and 6g, we show an application of our approach to a real-world problem – sorting fruits based on their level of ripeness. As ripeness is correlated with elasticity in climacteric fruits, by estimating the elasticity of each fruit during its grasp, the robot can sort them based on their ripeness in real time.

7 Limitations and Directions for Future Work

We have taken a first step towards robot grasping that minimizes the force used to grasp an object. In our work, we focused on the task of controlling force given a suitable grasp pose for a parallel-jaw gripper. To minimize the potential for damage, we would also need to address the problem of choosing grasp poses that do not require large forces (*e.g*., a banana picked up by one end). While we have used a parallel-jaw gripper, further analysis of the physics of grasping would be needed to extend our work to more complex end-effectors, such as dextrous hands that can cradle objects. Additionally, the performance of our method is tied to that of the tactile sensor we use. Current sensors trade off spatial and temporal resolution. While we addressed this problem in our work by using interpolation, to improve performance we will need sensors with both high update rate and spatial resolution. Furthermore, we have focused on the domain of fruits and vegetables, for which we could measure class priors. For objects without class priors, our method may still be effective, but this would need to be validated via extensive experiments. Finally, we have assumed that the object's material properties are more or less constant over its surface. We plan to extend our work to the harder problem of grasping objects with spatially varying properties (*e.g*., a stuffed animal).

8 Summary

We have introduced material-aware robotic grasping: a visuotactile approach that estimates instance-level material properties, online, during contact with an object. This allows us to drive closed-loop force control towards the theoretical minimum grasp force. We created *Squash*, a diverse dataset of 3D object models with measured material properties. Our visuotactile model, trained using *Squash*, recovers elasticity, mass, and friction parameters from tactile feedback. In simulation and on real hardware, our method consistently achieved gentle, non-damaging grasps. In addition, the *in situ* material estimation of our method opens avenues for object sorting, grasp-pose optimization, dexterous manipulation, or any task where knowing *how* an object deforms is of value.

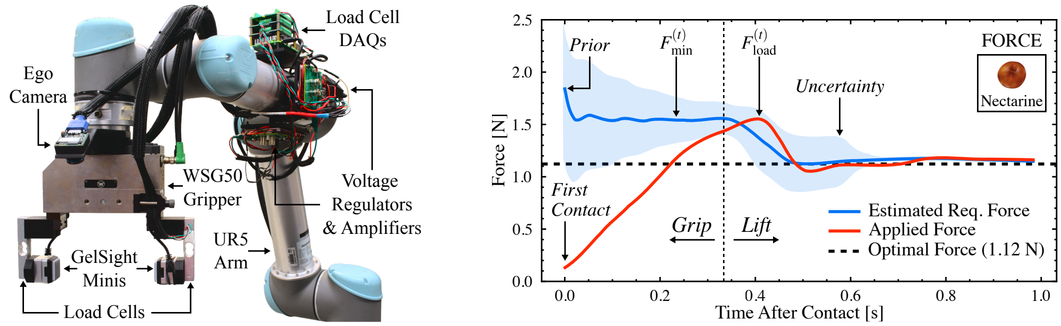

(a) Robot Arm and Gripper Hardware (b) Force Evolution Profile During a Grasp

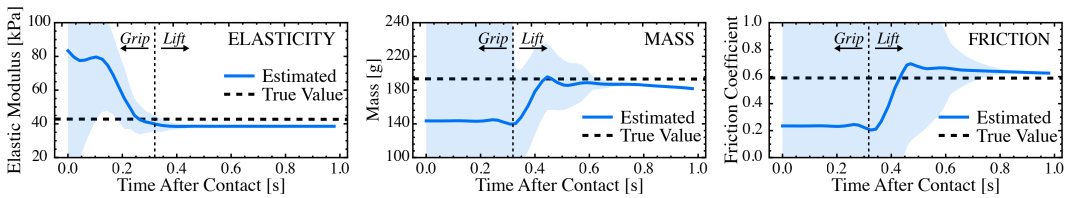

(c) Material Properties Estimation During a Grasp

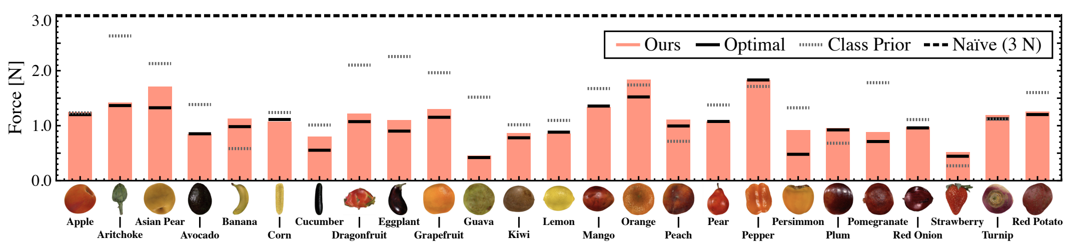

(d) Individual Predicted Grasp Forces

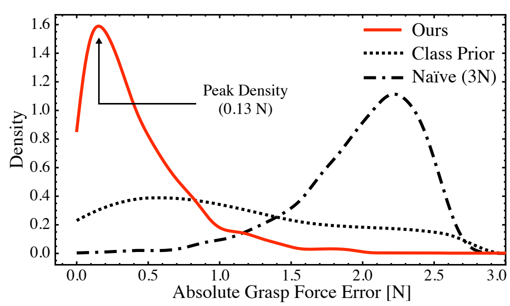

Approach	MAPE [%]	MAE [N]	Successful
Naïve (3N)	244.0±148.7	1.99±0.33	**29/29**
Class Prior	56.7±60.7	0.46±0.40	21/29
Ours	**23.0±17.9**	**0.21±0.14**	28/29

(e) Average Force Estimation Error

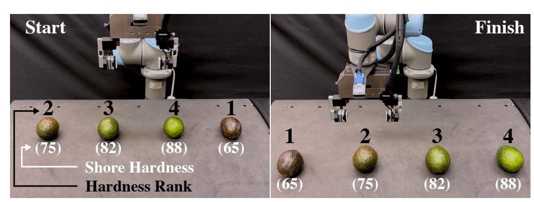

(f) Sorting *Avocados* Based on Ripeness

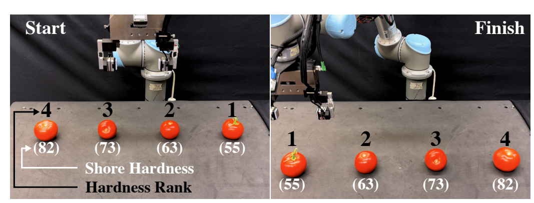

(g) Sorting *Tomatoes* Based on Ripeness

Figure 6. **Experimental Results.** (a) Robot hardware setup for real-world grasping experiments. Our approach is powered by GelSight Mini tactile sensors paired with bending beam load cells, but is not limited to this configuration. (b-c) Evolution of the applied grasp force and estimated material properties during a trial. The values converge quickly, and, as a result, we gently grasp the object in real time with a force close to the optimal force. (d) The applied grasp force by our approach for various real fruits and vegetables. We consistently apply at or slightly above the optimal force. (e) Performance of our approach over all 29 tested real fruits and vegetables, where we show accurate force control while maintaining a high grasp success rate. In the kernel density plot, our approach shows a tighter concentration around zero, indicating consistent performance that adapts to individual object properties. (f-g) Sorting fruits in real time based on their estimated elasticity, which is correlated with ripeness. This is one potential application of material-aware grasping beyond force control.

Acknowledgments

This work was supported in part by the Office of Naval Research (ONR) award N00014-24-1-2155 and in part by the National Science Foundation (NSF) and Center for Smart Streetscapes (CS3) under NSF Cooperative Agreement No. EEC-2133516. The authors are grateful to Behzad Kamgar-Parsi at ONR for his support. The authors thank Dhruv Yalamanchi and Isabel Tu for help with collecting the *Squash* dataset, Joaquin Palacios for discussions on robot control, and Silvia Sellán for visualization support.

References

[1] V. C. Paulun, U. Kleinholdermann, K. R. Gegenfurtner, J. B. J. Smeets, and E. Brenner. Center or side: Biases in selecting grasp points on small bars. *Experimental Brain Research*, 232(7): 2061–2072, July 2014. ISSN 1432-1106. doi:10.1007/s00221-014-3895-z.

[2] V. C. Paulun, K. R. Gegenfurtner, M. A. Goodale, and R. W. Fleming. Effects of material properties and object orientation on precision grip kinematics. *Experimental Brain Research*, 234(8):2253–2265, Aug. 2016. ISSN 1432-1106. doi:10.1007/s00221-016-4631-7.

[3] R. W. Fleming, C. Wiebel, and K. Gegenfurtner. Perceptual qualities and material classes. *Journal of Vision*, 13(8):9, July 2013. ISSN 1534-7362. doi:10.1167/13.8.9.

[4] R. W. Fleming. Visual perception of materials and their properties. *Vision Research*, 94:62–75, Jan. 2014. ISSN 0042-6989. doi:10.1016/j.visres.2013.11.004.

[5] R. W. Fleming. Material Perception. *Annual Review of Vision Science*, 3(Volume 3, 2017):365–388, Sept. 2017. ISSN 2374-4642, 2374-4650. doi:10.1146/annurev-vision-102016-061429.

[6] A. C. Zoeller, A. Lezkan, V. C. Paulun, R. W. Fleming, and K. Drewing. Integration of prior knowledge during haptic exploration depends on information type. *Journal of Vision*, 19(4): 20, Apr. 2019. ISSN 1534-7362. doi:10.1167/19.4.20.

[7] G. Maiello, M. Schepko, L. K. Klein, V. C. Paulun, and R. W. Fleming. Humans Can Visually Judge Grasp Quality and Refine Their Judgments Through Visual and Haptic Feedback. *Frontiers in Neuroscience*, 14, Jan. 2021. ISSN 1662-453X. doi:10.3389/fnins.2020.591898.

[8] M. Li, L. Zhang, T. Li, and Y. Jiang. Learning Gentle Grasping from Human-Free Force Control Demonstration. *IEEE Robotics and Automation Letters*, 10(3):2391–2398, Mar. 2025. ISSN 2377-3766, 2377-3774. doi:10.1109/LRA.2025.3530354.

[9] K. Nakahara and R. Calandra. Learning Gentle Grasping Using Vision, Sound, and Touch. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7607–7614, Oct. 2025. doi:10.1109/IROS60139.2025.11246722.

[10] J. Xu, M. Danielczuk, J. Ichnowski, J. Mahler, E. Steinbach, and K. Goldberg. Minimal Work: A Grasp Quality Metric for Deformable Hollow Objects. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1546–1552, May 2020. doi:10.1109/ICRA40945.2020.9197062.

[11] I. Huang, Y. Narang, R. Bajcsy, F. Ramos, T. Hermans, and D. Fox. DefGraspNets: Grasp Planning on 3D Fields with Graph Neural Nets. https://arxiv.org/abs/2303.16138v1, Mar. 2023.

[12] H. Zhang, K. Dana, and K. Nishino. Reflectance Hashing for Material Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3071–3080, 2015.

[13] K. J. Dana, B. Van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, Jan. 1999. ISSN 0730-0301, 1557-7368. doi:10.1145/300776.300778.

[14] M. Weinmann, J. Gall, and R. Klein. Material Classification Based on Training Data Synthesized Using a BTF Database. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8691, pages 156–171. Springer International Publishing, Cham, 2014. ISBN 978-3-319-10577-2 978-3-319-10578-9. doi:10.1007/978-3-319-10578-9_11.

[15] E. H. Adelson. On seeing stuff: The perception of materials by humans and machines. In *Human Vision and Electronic Imaging VI*, volume 4299, pages 1–12. SPIE, June 2001. doi:10.1117/12.429489.

[16] L. Sharan, C. Liu, R. Rosenholtz, and E. H. Adelson. Recognizing Materials Using Perceptually Inspired Features. *International Journal of Computer Vision*, 103(3):348–371, July 2013. ISSN 1573-1405. doi:10.1007/s11263-013-0609-0.

[17] S. Bell, P. Upchurch, N. Snavely, and K. Bala. Material Recognition in the Wild With the Materials in Context Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3479–3487, 2015.

[18] K. J. Dana. *Computational Texture and Patterns: From Textons to Deep Learning*. Synthesis Lectures on Computer Vision. Springer International Publishing, Cham, 2018. ISBN 978-3-031-00695-1 978-3-031-01823-7. doi:10.1007/978-3-031-01823-7.

[19] L. Liu, J. Chen, P. Fieguth, G. Zhao, R. Chellappa, and M. Pietikäinen. From BoW to CNN: Two Decades of Texture Representation for Texture Classification. *International Journal of Computer Vision*, 127(1):74–109, Jan. 2019. ISSN 1573-1405. doi:10.1007/s11263-018-1125-z.

[20] G. Schwartz and K. Nishino. Recognizing Material Properties from Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(8):1981–1995, Aug. 2020. ISSN 1939-3539. doi:10.1109/TPAMI.2019.2907850.

[21] M. Purri and K. Dana. Teaching Cameras to Feel: Estimating Tactile Physical Properties of Surfaces from Images. In A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, editors, *European Conference on Computer Vision*, volume 12372, pages 1–20. Springer International Publishing, Cham, 2020. ISBN 978-3-030-58582-2 978-3-030-58583-9. doi:10.1007/978-3-030-58583-9_1.

[22] M. Beveridge and S. K. Nayar. Hierarchical Material Recognition from Local Appearance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8165–8176, 2025.

[23] R. L. Klatzky, D. K. Pai, and E. P. Krotkov. Perception of Material from Contact Sounds. *Presence: Teleoperators and Virtual Environments*, 9(4):399–410, Aug. 2000. ISSN 1054-7460. doi:10.1162/105474600566907.

[24] C. Fang, D. Wang, F. Guo, J. Zou, and D. Song. A Fingertip Sensor and Algorithms for Pre-touch Distance Ranging and Material Detection in Robotic Grasping, Nov. 2023.

[25] T. Zhang, M. Sheinin, D. Chan, M. Rau, M. O'Toole, and S. G. Narasimhan. Analyzing Physical Impacts Using Transient Surface Wave Imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4339–4348, 2023.

[26] M. K. Johnson and E. H. Adelson. Retrographic sensing for the measurement of surface texture and shape. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1070–1077, June 2009. doi:10.1109/CVPR.2009.5206534.

[27] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson. Microgeometry capture using an elastomeric sensor. *ACM Transactions on Graphics (TOG)*, 30(4):46:1–46:8, July 2011. ISSN 0730-0301. doi:10.1145/2010324.1964941.

[28] W. Yuan, S. Dong, and E. H. Adelson. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors*, 17(12):2762, Dec. 2017. ISSN 1424-8220. doi:10.3390/s17122762.

[29] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, D. Jayaraman, and R. Calandra. DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor With Application to In-Hand Manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, July 2020. ISSN 2377-3766. doi:10.1109/LRA.2020.2977257.

[30] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson. Shape-independent hardness estimation using deep learning and a GelSight tactile sensor. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 951–958, May 2017. doi:10.1109/ICRA.2017.7989116.

[31] W. Yuan, R. Li, M. A. Srinivasan, and E. H. Adelson. Measurement of shear and slip with a GelSight tactile sensor. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 304–311, Seattle, WA, USA, May 2015. IEEE. ISBN 978-1-4799-6923-4. doi:10.1109/ICRA.2015.7139016.

[32] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta. ReSkin: Versatile, replaceable, lasting tactile skins. In *Proceedings of the 5th Conference on Robot Learning*, pages 587–597. PMLR, Jan. 2022.

[33] B. Huang and Y. Li. FlexiTac: A Low-Cost, Open-Source, Scalable Tactile Sensing Solution for Robotic Systems, Apr. 2026.

[34] V. Pattabiraman, Z. Huang, D. Panozzo, D. Zorin, L. Pinto, and R. Bhirangi. eFlesh: Highly customizable Magnetic Touch Sensing using Cut-Cell Microstructures, June 2025.

[35] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li. 3D-ViTac: Learning Fine-Grained Manipulation with Visuo-Tactile Sensing. https://arxiv.org/abs/2410.24091v2, Oct. 2024.

[36] X. Zhu, B. Huang, and Y. Li. Touch in the Wild: Learning Fine-Grained Manipulation with a Portable Visuo-Tactile Gripper. https://arxiv.org/abs/2507.15062v2, July 2025.

[37] W. Yuan, S. Wang, S. Dong, and E. Adelson. Connecting Look and Feel: Associating the Visual and Tactile Properties of Physical Materials. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5580–5588, 2017.

[38] Y. Li, J.-Y. Zhu, R. Tedrake, and A. Torralba. Connecting Touch and Vision via Cross-Modal Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10609–10618, 2019.

[39] X. Yu, R. Huang, C. Zhao, L. Zhou, and L. Ou. Def-Grasp: A Robot Grasping Detection Method for Deformable Objects Without Force Sensor. *Neural Processing Letters*, 55(8): 11739–11756, Dec. 2023. ISSN 1573-773X. doi:10.1007/s11063-023-11398-8.

[40] W. Ruan, W. Zhu, K. Wang, Q. Lu, W. Yeh, L. Luo, C. Su, and Q. Wang. Vision-Tactile Fusion Based Detection of Deformation and Slippage of Deformable Objects During Grasping. In F. Sun, A. Cangelosi, J. Zhang, Y. Yu, H. Liu, and B. Fang, editors, *Cognitive Systems and Information Processing*, pages 593–604, Singapore, 2023. Springer Nature. ISBN 978-981-9906-17-8. doi:10.1007/978-981-99-0617-8_43.

[41] N. Hogan. Impedance Control: An Approach to Manipulation: Part II—Implementation. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):8–16, Mar. 1985. ISSN 0022-0434. doi:10.1115/1.3140713.

[42] T. Li, C. Chi, C. Wang, N. Xue, C. Liu, and C. Liu. Real-Time and Autonomous Grasping Operation of Manipulator Based on Tactile Sensor Array. In *2018 2nd IEEE Advanced Information Management,Communicates,Electronic and Automation Control Conference (IMCEC)*, pages 732–736, May 2018. doi:10.1109/IMCEC.2018.8469277.

[43] S. J. Dharbaneshwer and A. Thondiyath. Contact Area-Based Modeling of Robotic Grasps Using Deformable Solid Mechanics. *International Journal of Applied Mechanics*, 13(03): 2150038, Apr. 2021. ISSN 1758-8251. doi:10.1142/S1758825121500381.

[44] S. D'Avella, M. Fontana, R. Vertechy, and P. Tripicchio. Towards autonomous soft grasping of deformable objects using flexible thin-film electro-adhesive gripper. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*, pages 1309–1314, Aug. 2022. doi:10.1109/CASE49997.2022.9926531.

[45] J. W. James and N. F. Lepora. Slip Detection for Grasp Stabilization With a Multifingered Tactile Robot Hand. *IEEE Transactions on Robotics*, 37(2):506–519, Apr. 2021. ISSN 1941-0468. doi:10.1109/TRO.2020.3031245.

[46] E. Judd, B. Aksoy, K. M. Digumarti, H. Shea, and D. Floreano. Slip Anticipation for Grasping Deformable Objects Using a Soft Force Sensor. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10003–10008, Oct. 2022. doi:10.1109/IROS47612.2022.9981174.

[47] T. N. Le, J. Lundell, and V. Kyrki. Safe Grasping with a Force Controlled Soft Robotic Hand. In *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 342–349, Oct. 2020. doi:10.1109/SMC42975.2020.9283496.

[48] H. Dong, C.-Y. Chen, C. Qiu, C.-H. Yeow, and H. Yu. GSG: A Granary Soft Gripper with Mechanical Force Sensing via 3-Dimensional Snap-Through Structure, Nov. 2021.

[49] T. N. Le, J. Lundell, F. J. Abu-Dakka, and V. Kyrki. Towards synthesizing grasps for 3D deformable objects with physics-based simulation. https://arxiv.org/abs/2107.08898v1, July 2021.

[50] T. N. Le, J. Lundell, F. J. Abu-Dakka, and V. Kyrki. Deformation-Aware Data-Driven Grasp Synthesis. *IEEE Robotics and Automation Letters*, 7(2):3038–3045, Apr. 2022. ISSN 2377-3766. doi:10.1109/LRA.2022.3146551.

[51] X. Lin, C. Qi, Y. Zhang, Z. Huang, K. Fragkiadaki, Y. Li, C. Gan, and D. Held. Planning with Spatial-Temporal Abstraction from Point Clouds for Deformable Object Manipulation. In *Proceedings of The 6th Conference on Robot Learning*, pages 1640–1651. PMLR, Mar. 2023.

[52] K. Zhang, B. Li, K. Hauser, and Y. Li. AdaptiGraph: Material-Adaptive Graph-Based Neural Dynamics for Robotic Manipulation, July 2024.

[53] K. Zhang, B. Li, K. Hauser, and Y. Li. Particle-Grid Neural Dynamics for Learning Deformable Object Models from RGB-D Videos, Nov. 2025.

[54] H. Jiang, H.-Y. Hsu, K. Zhang, H.-N. Yu, S. Wang, and Y. Li. PhysTwin: Physics-Informed Reconstruction and Simulation of Deformable Objects from Videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7219–7230, 2025.

[55] H. Zhang, G. Zheng, K. Zhang, J. Song, S. Patel, S. Hu, Y. Li, C. Zheng, and P. Y. Chen. BoxTwin: Learning Elastoplastic Articulated Object Dynamics from Videos.

[56] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine. More Than a Feeling: Learning to Grasp and Regrasp using Vision and Touch. *IEEE Robotics and Automation Letters*, 3(4):3300–3307, Oct. 2018. ISSN 2377-3766, 2377-3774. doi:10.1109/LRA.2018.2852779.

[57] S. Cui, R. Wang, J. Wei, F. Li, and S. Wang. Grasp State Assessment of Deformable Objects Using Visual-Tactile Fusion Perception. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 538–544, May 2020. doi:10.1109/ICRA40945.2020.9196787.

[58] Y. Luo, C. Liu, Y. J. Lee, J. DelPreto, K. Wu, M. Foshey, D. Rus, T. Palacios, Y. Li, A. Torralba, and W. Matusik. Adaptive tactile interaction transfer via digitally embroidered smart gloves. *Nature Communications*, 15(1):868, Jan. 2024. ISSN 2041-1723. doi:10.1038/s41467-024-45059-8.

[59] Y. Wi, J. Yin, E. Xiang, A. Sharma, J. Malik, M. Mukadam, N. Fazeli, and T. Hellebrekers. TactAlign: Human-to-Robot Policy Transfer via Tactile Alignment. https://arxiv.org/abs/2602.13579v1, Feb. 2026.

[60] E. Helmut, N. Funk, T. Schneider, C. de Farias, and J. Peters. Tactile-Conditioned Diffusion Policy for Force-Aware Robotic Manipulation. https://arxiv.org/abs/2510.13324v1, Oct. 2025.

[61] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in Manipulation Research: The YCB Object and Model Set and Benchmarking Protocols. https://arxiv.org/abs/1502.03143v1, Feb. 2015.

[62] H. Liang, X. Ma, S. Li, M. Görner, S. Tang, B. Fang, F. Sun, and J. Zhang. PointNetGPD: Detecting Grasp Configurations from Point Sets. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3629–3635, May 2019. doi:10.1109/ICRA.2019.8794435.

[63] A. Mousavian, C. Eppner, and D. Fox. 6-DOF GraspNet: Variational Grasp Generation for Object Manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2901–2910, 2019.

[64] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee. Learning 6-DOF Grasping Interaction via Deep Geometry-Aware 3D Representations. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3766–3773, May 2018. doi:10.1109/ICRA.2018.8460609.

[65] H.-S. Fang, C. Wang, M. Gou, and C. Lu. GraspNet-1Billion: A Large-Scale Benchmark for General Object Grasping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11444–11453, 2020.

[66] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. https://arxiv.org/abs/1512.03012v1, Dec. 2015.

[67] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke. Google Scanned Objects: A High-Quality Dataset of 3D Scanned Household Items. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2553–2560, May 2022. doi:10.1109/ICRA46639.2022.9811809.

[68] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi. Objaverse: A Universe of Annotated 3D Objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023.

[69] J. Cao and E. Kalogerakis. SOPHY: Learning to Generate Simulation-Ready Objects with Physical Materials. https://arxiv.org/abs/2504.12684v3, Apr. 2025.

[70] I. Huang, Y. Narang, C. Eppner, B. Sundaralingam, M. Macklin, R. Bajcsy, T. Hermans, and D. Fox. DefGraspSim: Physics-Based Simulation of Grasp Outcomes for 3D Deformable Objects. *IEEE Robotics and Automation Letters*, 7(3):6274–6281, July 2022. ISSN 2377-3766. doi:10.1109/LRA.2022.3158725.

[71] F. P. Bowden and D. Tabor. *The Friction and Lubrication of Solids*. Clarendon Press, 2001. ISBN 978-0-19-850777-2.

[72] D. A. H. Hanaor, Y. Gan, and I. Einav. Static friction at fractal interfaces. *Tribology International*, 93:229–238, Jan. 2016. ISSN 0301-679X. doi:10.1016/j.triboint.2015.09.016.

[73] H. Hertz. On the contact of elastic solids. In *Miscellaneous Papers*, chapter 5, pages 146–162. Macmillan, London, 1896.

[74] H. Hertz. On the contact of rigid elastic solids and on hardness. In *Miscellaneous Papers*, chapter 6, pages 163–183. Macmillan, London, 1896.

[75] A. N. Gent. On the Relation between Indentation Hardness and Young's Modulus. *Rubber Chemistry and Technology*, 31(4):896–906, Sept. 1958. ISSN 1943-4804, 0035-9475. doi:10.5254/1.3542351.

[76] A. W. Mix and A. J. Giacomin. Standardized Polymer Durometry. *Journal of Testing and Evaluation*, 39(4):696–705, July 2011. ISSN 0090-3973. doi:10.1520/JTE103205.

[77] 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*, 39(4), Aug. 2020. ISSN 0730-0301, 1557-7368. doi:10.1145/3386569.3392425.

[78] K. Huang, F. Chitalu, H. Lin, and T. Komura. GIPC: Fast and stable Gauss-Newton optimization of IPC barrier energy. *ACM Transactions on Graphics*, 43(2):1–18, Apr. 2024. ISSN 0730-0301, 1557-7368. doi:10.1145/3643028.

[79] K. Huang, X. Lu, H. Lin, T. Komura, and M. Li. StiffGIPC: Advancing GPU IPC for Stiff Affine-Deformable Simulation. *ACM Transactions on Graphics*, 44(3):31:1–31:20, May 2025. ISSN 0730-0301. doi:10.1145/3735126.

[80] C. M. Kim, M. Danielczuk, I. Huang, and K. Goldberg. IPC-GraspSim: Reducing the Sim2Real Gap for Parallel-Jaw Grasping with the Incremental Potential Contact Model, Mar. 2022.

[81] D. H. Nguyen, T. Schneider, G. Duret, A. Kshirsagar, B. Belousov, and J. Peters. TacEx: GelSight Tactile Simulation in Isaac Sim – Combining Soft-Body and Visuotactile Simulators, Nov. 2024.

[82] R. Hasani, M. Lechner, A. Amini, L. Liebenwein, A. Ray, M. Tschaikowski, G. Teschl, and D. Rus. Closed-form continuous-time neural networks. *Nature Machine Intelligence*, 4(11): 992–1003, Nov. 2022. ISSN 2522-5839. doi:10.1038/s42256-022-00556-7.

[83] C. Ferrari and J. Canny. Planning optimal grasps. In *Proceedings 1992 IEEE International Conference on Robotics and Automation*, pages 2290–2295 vol.3, May 1992. doi:10.1109/ROBOT.1992.219918.

[84] A. Miller and P. Allen. Graspit! A versatile simulator for robotic grasping. *IEEE Robotics & Automation Magazine*, 11(4):110–122, Dec. 2004. ISSN 1558-223X. doi:10.1109/MRA.2004.1371616.