

Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics

Jeffrey Mahler*, Jacky Liang*, Sherdil Niyaz*, Michael Laskey*, Richard Doan*, Xinyu Liu*,
Juan Aparicio Ojea†, and Ken Goldberg*

*Dept. of EECS, University of California, Berkeley

Email: {jmahler, jackyliang, sniyaz, laskeymd, rdoan, xinyuliu, goldberg}@berkeley.edu

† Siemens Corporation, Corporate Technology

Email: juan.aparicio@siemens.com

Abstract—To reduce data collection time for deep learning of robust robotic grasp plans, we explore training from a synthetic dataset of 6.7 million point clouds, grasps, and analytic grasp metrics generated from thousands of 3D models from Dex-Net 1.0 in randomized poses on a table. We use the resulting dataset, Dex-Net 2.0, to train a Grasp Quality Convolutional Neural Network (GQ-CNN) model that rapidly predicts the probability of success of grasps from depth images, where grasps are specified as the planar position, angle, and depth of a gripper relative to an RGB-D sensor. Experiments with over 1,000 trials on an ABB YuMi comparing grasp planning methods on singulated objects suggest that a GQ-CNN trained with only synthetic data from Dex-Net 2.0 can be used to plan grasps in 0.8s with a success rate of 93% on eight known objects with adversarial geometry and is 3× faster than registering point clouds to a precomputed dataset of objects and indexing grasps. The Dex-Net 2.0 grasp planner also has the highest success rate on a dataset of 10 novel rigid objects and achieves 99% precision (one false positive out of 69 grasps classified as robust) on a dataset of 40 novel household objects, some of which are articulated or deformable. Code, datasets, videos, and supplementary material are available at <http://berkeleyautomation.github.io/dex-net>.

I. INTRODUCTION

Reliable robotic grasping is challenging due to imprecision in sensing and actuation, which leads to uncertainty about properties such as object shape, pose, material properties, and mass. Recent results suggest that deep neural networks trained on large datasets of human grasp labels [32] or physical grasp outcomes [41] can be used to plan grasps that are successful across a wide variety of objects directly from images or point clouds, similar to generalization results in computer vision [29]. However, data collection requires either tedious human labeling [26] or months of execution time on a physical system [33].

An alternative approach is to plan grasps using physics-based analyses such as caging [47], grasp wrench space (GWS) analysis [45], robust GWS analysis [56], or simulation [26], which can be rapidly computed using Cloud Computing [28]. However, these methods assume a separate perception system that estimates properties such as object shape or pose either perfectly [45] or according to known Gaussian distributions [35]. This is prone to errors [2], may




Fig. 1: Dex-Net 2.0 Architecture. (Center) The Grasp Quality Convolutional Neural Network (GQ-CNN) is trained offline to predict the robustness candidate grasps from depth images using a dataset of 6.7 million synthetic point clouds, grasps, and associated robust grasp metrics computed with Dex-Net 1.0. (Left) When an object is presented to the robot, a depth camera returns a 3D point cloud, where pairs of antipodal points identify a set of several hundred grasp candidates. (Right) The GQ-CNN rapidly determines the most robust grasp candidate, which is executed with the ABB YuMi robot.

not generalize well to new objects, and can be slow to match point clouds to known models during execution [14]. In this paper we instead consider predicting grasp success directly from depth images by training a deep Convolutional Neural Network (CNN) on a massive dataset of parallel-jaw grasps, grasp metrics, and rendered point clouds generated using analytic models of robust grasping and image formation [19, 36], building upon recent research on classifying force closure grasps [16, 51] and the outcomes of dynamic grasping simulations [25, 26, 54].

Our primary contributions are: 1) the Dexterity Network (Dex-Net) 2.0, a dataset associating 6.7 million point clouds and analytic grasp quality metrics with parallel-jaw grasps planned using robust quasi-static GWS analysis on a dataset of 1,500 3D object models, 2) a Grasp Quality Convolutional Neural Network (GQ-CNN) model trained to classify robust grasps in depth images using expected epsilon quality as supervision, where each grasp is specified as a planar pose and depth relative to a camera, and 3) a grasp planning method that samples antipodal grasp candidates and ranks them with a GQ-CNN.

In over 1,000 physical trials of grasping single objects on a tabletop with an ABB YuMi robot, we compare Dex-Net 2.0 to image-based grasp heuristics, a random forest [51], an

SVM [52], and a baseline that recognizes objects, registers their 3D pose [14], and indexes Dex-Net 1.0 [35] for the most robust grasp to execute. We find that the Dex-Net 2.0 grasp planner is $3\times$ faster than the registration-based method, 93% successful on objects seen in training (the highest of learning-based methods), and is the best performing method on novel objects, achieving 99% precision on a dataset of 40 household objects despite being trained entirely on synthetic data.

II. RELATED WORK

Grasp Planning. Given an object and reachability constraints due to the environment, grasp planning considers finding a gripper configuration that maximizes a success (or quality) metric. Methods fall into one of two categories based on success criteria: *analytic* methods [45], which consider performance according to physical models such as the ability to resist external wrenches [43], and *empirical* (or data-driven) methods [4], which typically use human labels [2] or the ability to lift the object in physical trials [41].

Analytic Methods. Analytic approaches typically assume that object and contact locations are known exactly and consider either the ability to resist external wrenches [45] or the ability to constrain the object’s motion [47]. To execute grasps on a physical robot, a common approach is to precompute a database of known 3D objects labeled with grasps and quality metrics such as GraspIt! [15]. Precomputed grasps are indexed using *point cloud registration*: matching point clouds to known 3D object models in the database using visual and geometric similarity [4, 5, 8, 14, 21, 23, 28] and executing the highest quality grasp for the estimated object instance.

Robust grasp planning (RGP) methods maximize grasp *robustness*, or the expected value of an analytic metric under uncertainty in sensing and control. This typically involves labeling grasps on a database 3D object models with robust metrics such as probability of force closure [28] or the pose error robust metric [56] and using registration-based planning. Recent research has demonstrated that the sampling complexity of RGP can be improved using Multi-Armed Bandits [30] and datasets of prior 3D objects and robust grasps, such as the Dexterity Network (Dex-Net) 1.0 [35]. In this work we make a major extension to Dex-Net 1.0 by associating synthetic point clouds with robust grasps and training a Convolutional Neural Network to map point clouds and candidates grasps to estimated robustness.

Empirical Methods. Empirical approaches typically use machine learning to develop models that map from robotic sensor readings directly to success labels from humans or physical trials. Human labels have become popular due to empirical correlation with physical success [2], although they may be expensive to acquire for large datasets. Research in this area has largely focused on associating human labels with graspable regions in RGB-D images [32] or point clouds [10, 22, 26]. Lenz et al. [32] created a dataset of over 1k RGB-D images with human labels of successful and unsuccessful grasping regions, which has been used to train fast CNN-based detection models [46].

Another line of research on empirical grasp planning has attempted to optimize success in physical trials directly. The time cost of generating samples on a physical robot led to the development active methods for acquiring grasping experiences such as Multi-Armed Bandits using Correlated Beta Processes [38] or Prior Confidence Bounds [40]. Recently Pinto and Gupta [41] scaled up data collection by recording over 40k grasping experiences on a Baxter and iteratively training CNNs to predict lifting successes or to resist grasp perturbations caused by an adversary [42]. Levine et al. [33] scaled up dataset collection even further, collecting over 800k datapoints with a set of continuously running robotic arms and using deep learning to predict end effector poses. However, this required over 2 months of training across up to 14 robots.

Computer Vision Techniques in Robot Grasping. Grasps for a physical robot are typically planned from images of target objects. Analytic grasp planning methods register images of rigid objects to a known database of 3D models, which typically involves segmentation, classification, and geometric pose estimation [8, 23, 49, 58] from 3D point cloud data in order to index precomputed grasps. However, this multi-stage approach can be prone to many hyperparameters that are difficult to tune and errors can compound across modules.

An alternative approach is to use deep learning to estimate 3D object shape and pose directly from color and depth images [18, 59]. Recent research in robotics has focused on how to improve accuracy in object recognition by structuring the way the neural network fuses the separate color and depth streams from images [44] and adding synthetic noise to synthetic training images [11]. Another approach is to detect graspable regions directly in images without explicitly representing object shape and pose [32, 39, 41, 50], as it may not always be necessary to explicitly recognize objects and their pose to perform a grasp.

Since training on real images may require significant data collection time, an alternative approach is to learn on simulated images [54, 55] and to adapt the representation to real data [53, 60]. Recent research suggests that in some cases it may be sufficient to train on datasets generated using perturbations to the parameters of the simulator [16, 48]. Notably, Johns et al. [25] used rendered depth images with simulated noise to train a CNN-based detector for planar grasps using dynamic simulations as supervision. We build upon these results by using robust analytic grasp metrics as supervision, using the gripper’s distance from the camera in predictions, and performing extensive evaluations on a physical robot.

III. PROBLEM STATEMENT

We consider the problem of planning a robust planar parallel-jaw grasp for a singulated rigid object resting on a table based on point clouds from a depth camera. We learn a function that takes as input a candidate grasp and a depth image and outputs an estimate of robustness [28, 56], or probability of success under uncertainty in sensing and control.

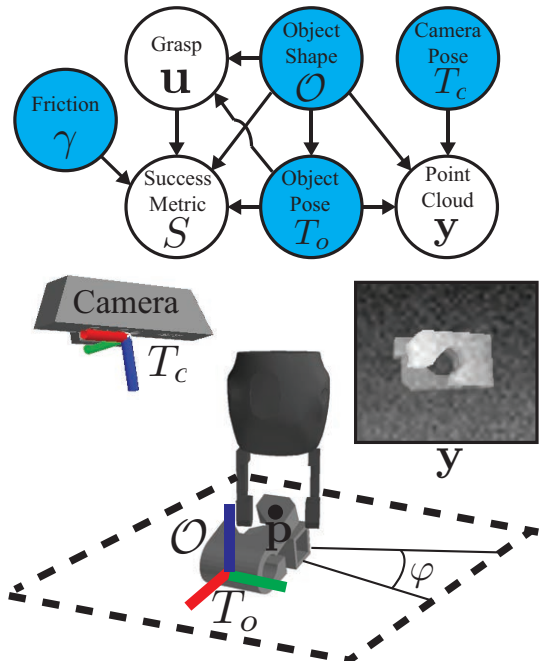


Fig. 2: Graphical model for robust parallel-jaw grasping of objects on a table surface based on point clouds. Blue nodes are variables included in the state representation. Object shapes \mathcal{O} are uniformly distributed over a discrete set of object models and object poses T_o are distributed over the object’s stable poses and a bounded region of a planar surface. Grasps $\mathbf{u} = (\mathbf{p}, \varphi)$ are sampled uniformly from the object surface using antipodality constraints. Given the coefficient of friction γ we evaluate an analytic success metric S for a grasp on an object. A synthetic 2.5D point cloud \mathbf{y} is generated from 3D meshes based on the camera pose T_c , object shape, and pose and corrupted with multiplicative and Gaussian Process noise.

A. Assumptions

We assume a parallel-jaw gripper, rigid objects singulated on a planar worksurface, and single-view (2.5D) point clouds taken with a depth camera. For generating datasets, we assume a known gripper geometry and a single overhead depth camera with known intrinsics.

B. Definitions

States. Let $\mathbf{x} = (\mathcal{O}, T_o, T_c, \gamma)$ denote a state describing the variable properties of the camera and objects in the environment, where \mathcal{O} specifies the geometry and mass properties of an object, T_o, T_c are the 3D poses of the object and camera, respectively, and $\gamma \in \mathbb{R}$ is the coefficient of friction between the object and gripper.

Grasps. Let $\mathbf{u} = (\mathbf{p}, \varphi) \in \mathbb{R}^3 \times \mathcal{S}^1$ denote a parallel-jaw grasp in 3D space specified by a center $\mathbf{p} = (x, y, z) \in \mathbb{R}^3$ relative to the camera and an angle in the table plane $\varphi \in \mathcal{S}^1$.

Point Clouds. Let $\mathbf{y} = \mathbb{R}_+^{H \times W}$ be a 2.5D point cloud represented as a depth image with height H and width W taken by a camera with known intrinsics [19], and let T_c be the 3D pose of the camera.

Robust Analytic Grasp Metrics. Let $S(\mathbf{u}, \mathbf{x}) \in \{0, 1\}$ be a binary-valued grasp success metric, such as force closure or physical lifting. Let $p(S, \mathbf{u}, \mathbf{x}, \mathbf{y})$ be a joint distribution on grasp success, grasps, states, and point clouds modeling imprecision in sensing and control. For example, p could be

defined by noisy sensor readings of a known set of industrial parts coming down a conveyor belt in arbitrary poses. Let the *robustness* of a grasp given an observation [5, 56] be the expected value of the metric, or probability of success under uncertainty in sensing and control: $Q(\mathbf{u}, \mathbf{y}) = \mathbb{E}[S | \mathbf{u}, \mathbf{y}]$.

C. Objective

Our goal is to learn a robustness function $Q_{\theta^*}(\mathbf{u}, \mathbf{y}) \in [0, 1]$ over many possible grasps, objects, and images that classifies grasps according to the binary success metric:

$$\theta^* = \operatorname{argmin}_{\theta \in \Theta} \mathbb{E}_{p(S, \mathbf{u}, \mathbf{x}, \mathbf{y})} [\mathcal{L}(S, Q_{\theta}(\mathbf{u}, \mathbf{y}))] \quad (\text{III.1})$$

where \mathcal{L} is the cross-entropy loss function and Θ defines the parameters of the Grasp Quality Convolutional Network (GQ-CNN) described in Section IV-B. This objective is motivated by that fact that $Q_{\theta^*} = Q$ for all possible grasps and images when there exists some $\theta \in \Theta$ such that $Q_{\theta} = Q$ [37]. The estimated robustness function can be used in a grasping policy that maximizes Q_{θ^*} over a set of candidate grasps: $\pi_{\theta}(\mathbf{y}) = \operatorname{argmax}_{\mathbf{u} \in \mathcal{C}} Q_{\theta}(\mathbf{u}, \mathbf{y})$, where \mathcal{C} specifies constraints on the set of available grasps, such as collisions or kinematic feasibility. Learning Q rather than directly learning the policy allows us to enforce task-specific constraints without having to update the learned model.

IV. LEARNING A GRASP ROBUSTNESS FUNCTION

Solving for the grasp robustness function in objective III.1 is challenging for several reasons. First, we may need a huge number of samples to approximate the expectation over a large number of possible objects. We address this by generating Dex-Net 2.0, a training dataset of 6.7 million synthetic point clouds, parallel-jaw grasps, and robust analytic grasp metrics across 1,500 3D models sampled from the graphical model in Fig. 2. Second, the relationship between point clouds, grasps, and metrics over a large dataset of objects may be complex and difficult to learn with linear or kernelized models. Consequently, we develop a Grasp Quality Convolutional Neural Network (GQ-CNN) model that classifies robust grasp poses in depth images and train the model on data from Dex-Net 2.0.

A. Dataset Generation

We estimate Q_{θ^*} using a sample approximation [12] of the objective in Equation III.1 using i.i.d samples $(S_1, \mathbf{u}_1, \mathbf{x}_1, \mathbf{y}_1), \dots, (S_N, \mathbf{u}_N, \mathbf{x}_N, \mathbf{y}_N) \sim p(S, \mathbf{u}, \mathbf{x}, \mathbf{y})$ from our generative graphical model for images, grasps, and success metrics:

$$\hat{\theta} = \operatorname{argmin}_{\theta \in \Theta} \sum_{i=1}^N \mathcal{L}(S_i, Q_{\theta}(\mathbf{u}_i, \mathbf{y}_i)).$$

1) *Graphical Model:* Our graphical model is illustrated in Fig. 2 and models $p(S, \mathbf{u}, \mathbf{x}, \mathbf{y})$ as the product of a state distribution $p(\mathbf{x})$, an observation model $p(\mathbf{y}|\mathbf{x})$, a grasp candidate model $p(\mathbf{u}|\mathbf{x})$, and an analytic model of grasp success $p(S|\mathbf{u}, \mathbf{x})$.

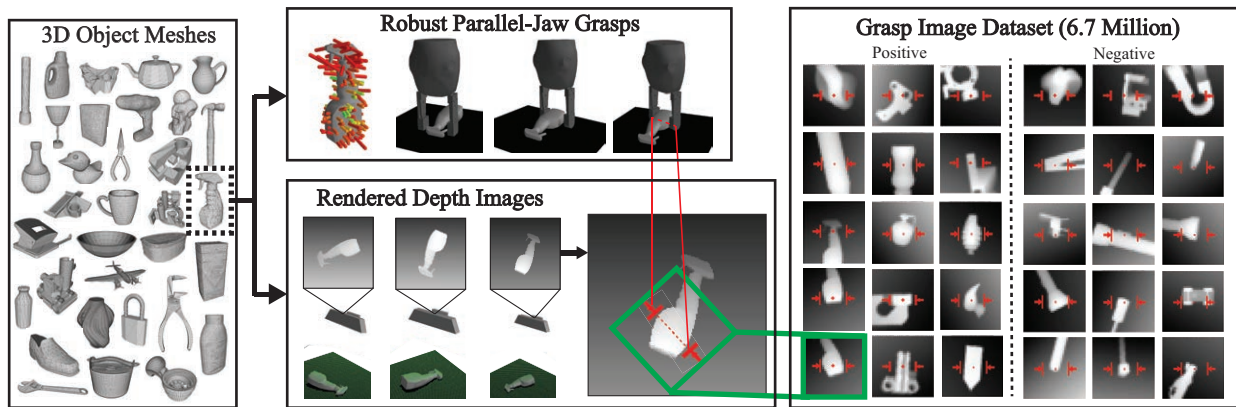


Fig. 3: Dex-Net 2.0 pipeline for training dataset generation. (Left) The database contains 1,500 3D object mesh models. (Top) For each object, we sample hundreds of parallel-jaw grasps to cover the surface and evaluate robust analytic grasp metrics using sampling. For each stable pose of the object we associate a set of grasps that are perpendicular to the table and collision-free for a given gripper model. (Bottom) We also render point clouds of each object in each stable pose, with the planar object pose and camera pose sampled uniformly at random. Every grasp for a given stable pose is associated with a pixel location and orientation in the rendered image. (Right) Each image is rotated, cropped, and scaled to align the grasp pixel location with the image center and the grasp axis with the middle row of the image, creating a 32×32 grasp image. The full dataset contains over 6.7 million grasp images.

Distribution	Description
$p(\gamma)$	truncated Gaussian distribution over friction coefficients
$p(\mathcal{O})$	discrete uniform distribution over 3D object models
$p(T_o \mathcal{O})$	continuous uniform distribution over the discrete set of object stable poses and planar poses on the table surface
$p(T_c)$	continuous uniform distribution over spherical coordinates for radial bounds $[r_\ell, r_u]$ and polar angle in $[0, \delta]$

TABLE I: Details of the distributions used in the Dex-Net 2.0 graphical model for generating the Dex-Net training dataset.

We model the state distribution as

$$p(\mathbf{x}) = p(\gamma)p(\mathcal{O})p(T_o|\mathcal{O})p(T_c)$$

where the distributions are detailed in Table I. Our grasp candidate model $p(\mathbf{u} | \mathbf{x})$ is a uniform distribution over pairs of antipodal contact points on the object surface that form a grasp axis parallel to the table plane. Our observation model is $\mathbf{y} = \alpha\hat{\mathbf{y}} + \epsilon$ where $\hat{\mathbf{y}}$ is a rendered depth image for a given object in a given pose, α is a Gamma random variable modeling depth-proportional noise, and ϵ is zero-mean Gaussian Process noise over pixel coordinates with bandwidth ℓ and measurement noise σ modeling additive noise [36]. We model grasp success as:

$$S(\mathbf{u}, \mathbf{x}) = \begin{cases} 1 & E_Q > \delta \text{ and } \text{collfree}(\mathbf{u}, \mathbf{x}) \\ 0 & \text{otherwise} \end{cases}$$

where E_Q is the robust epsilon quality defined in [51], a variant of the pose error robust metric [56] that includes uncertainty in friction and gripper pose, and $\text{collfree}(\mathbf{u}, \mathbf{x})$ indicates that the gripper does not collide with the object or table. The supplemental file details the parameters of these distributions.

2) *Database*: Dex-Net 2.0 contains 6.7 million datapoints generated using the pipeline of Fig. 3.

3D Models. The dataset contains a subset of 1,500 mesh models from Dex-Net 1.0: 1,371 synthetic models from

3DNet [57] and 129 laser scans from the KIT object database [27]. Each mesh is aligned to a standard frame of reference using the principal axes, rescaled to fit within a gripper width of 5.0cm (the opening width of an ABB YuMi gripper), and assigned a mass of 1.0kg centered in the object bounding box since some meshes are nonclosed. For each object we also compute a set of stable poses [13] and store all stable poses with probability of occurrence above a threshold.

Parallel-Jaw Grasps. Each object is labeled with a set of up to 100 parallel-jaw grasps. The grasps are sampled using the rejection sampling method for antipodal point pairs developed in Dex-Net 1.0 [35] with constraints to ensure coverage of the object surface [34]. For each grasp we evaluate the expected epsilon quality E_Q [43] under object pose, gripper pose, and friction coefficient uncertainty using Monte-Carlo sampling [51],

Rendered Point Clouds. Every object is also paired with a set of 2.5D point clouds (depth images) for each object stable pose, with camera poses and planar object poses sampled according to the graphical model described in Section IV-A1. Images are rendered using a pinhole camera model and perspective projection with known camera intrinsics, and each rendered image is centered on the object of interest using pixel transformations. Noise is added to the images during training as described in Section IV-B3.

B. Grasp Quality Convolutional Neural Network

1) *Architecture*: The Grasp Quality Convolutional Neural Network (GQ-CNN) architecture, illustrated in Fig. 4 and detailed in the caption, defines the set of parameters Θ used to represent the grasp robustness function Q_θ . The GQ-CNN takes as input the gripper depth from the camera z and a depth image centered on the grasp center pixel $\mathbf{v} = (i, j)$ and aligned with the grasp axis orientation φ . The image-gripper alignment removes the need to learn rotational invariances that can be modeled by known, computationally-efficient image




Fig. 4: (Left) Architecture of the Grasp Quality Convolutional Neural Network (GQ-CNN). Planar grasp candidates $\mathbf{u} = (i, j, \varphi, z)$ are generated from a depth image and transformed to align the image with the grasp center pixel (i, j) and orientation φ . The architecture contains four convolutional layers in pairs of two separated by ReLU nonlinearities followed by 3 fully connected layers and a separate input layer for the z , the distance of the gripper from the camera. The use of convolutional layers was motivated by the relevance of depth edges as features for learning in previous research [3, 32, 35] and the use of ReLUs was motivated by image classification results [29]. The network estimates the probability of grasp success (robustness) $Q_\theta \in [0, 1]$, which can be used to rank grasp candidates. (Right) The first layer of convolutional filters learned by the GQ-CNN on Dex-Net 2.0. The filters appear to compute oriented image gradients at various scales, which may be useful for inferring contact normals and collisions between the gripper and object.

transformations (similar to spatial transformer networks [24]) and allows the network to evaluate any grasp orientation in the image rather than a predefined discrete set as in [25, 41]. Following standard preprocessing conventions, we normalize the input data by subtracting the mean and dividing by the standard deviation of the training data and then pass the image and gripper depth through the network to estimate grasp robustness. The GQ-CNN has approximately 18 million parameters.

2) *Training Dataset*: GQ-CNN training datasets are generated by associating grasps with a pixel \mathbf{v} , orientation φ , and depth z relative to rendered depth images as illustrated in Fig. 3. We compute these parameters by transforming grasps into the camera frame of reference using the camera pose T_c and projecting the 3D grasp position and orientation onto the imaging plane of the camera [19]. We then transform all pairs of images and grasp configurations to a single image centered on \mathbf{v} and oriented along φ (see the left panel of Fig. 4 for an illustration). The Dex-Net 2.0 training dataset contains 6.7 million datapoints and approximately 21.2% positive examples for the thresholded robust epsilon quality with threshold $\delta = 0.002$ [26] and a custom YuMi gripper.

3) *Optimization*: We optimize the parameters of the GQ-CNN using backpropagation with stochastic gradient descent and momentum [29]. We initialize the weights of the model by sampling from a zero mean gaussian with variance $\frac{2}{n_i}$, where n_i is the number of inputs to the i -th network layer [20]. To augment the dataset, we reflect the image about its vertical and horizontal axes and rotate each image by 180° since these lead to equivalent grasps. We also adaptively sample image noise from our noise model (see Section IV-A1) before computing the batch gradient for new samples during training to model imaging noise without explicitly storing multiple versions of each image. To speed up noise sampling we approximate the Gaussian Process noise by upsampling an array of uncorrelated zero-mean Gaussian noise using bilinear interpolation. We set hyperparameters based on the performance on a randomize

synthetic validation set as described in Section VI-C.

V. GRASP PLANNING

The Dex-Net 2.0 grasp planner uses the robust grasping policy $\pi_\theta(\mathbf{y}) = \operatorname{argmax}_{\mathbf{u} \in \mathcal{C}} Q_\theta(\mathbf{u}, \mathbf{y})$ illustrated in Fig. 1. The set \mathcal{C} is a discrete set of antipodal candidate grasps [7] sampled uniformly at random in image space for surface normals defined by the depth image gradients. Each grasp candidate is evaluated by the GQ-CNN, and the most robust grasp that is (a) kinematically reachable and (b) not in collision with the table is executed. The supplemental file contains additional details.

VI. EXPERIMENTS

We evaluated classification performance on both real and synthetic data and performed extensive physical evaluations on an ABB YuMi with custom silicone gripper tips designed by Guo et al. [17] to benchmark the performance of grasping a single object. All experiments ran on a Desktop running Ubuntu 14.04 with a 2.7 GHz Intel Core i5-6400 Quad-Core CPU and an NVIDIA GeForce 980, and we used an NVIDIA GeForce GTX 1080 for training large models.

A. Physical Benchmark Description

We created a benchmark for grasping single objects on a tabletop to compare grasp planning methods. The setup is illustrated in Fig. 5 and the experimental procedure is described in the caption and shown in the supplemental video¹. Each grasp planner received as input a color image, depth image, bounding box containing the object, and camera intrinsics, and output a target grasping pose for the gripper. A human operator was required to reset the object in the workspace on each trial, and therefore blinded operators from which grasp planning method was being tested in order to remove bias.

We compared performance on this benchmark with the following metrics:

¹<https://youtu.be/9eqAxk95I3Y>

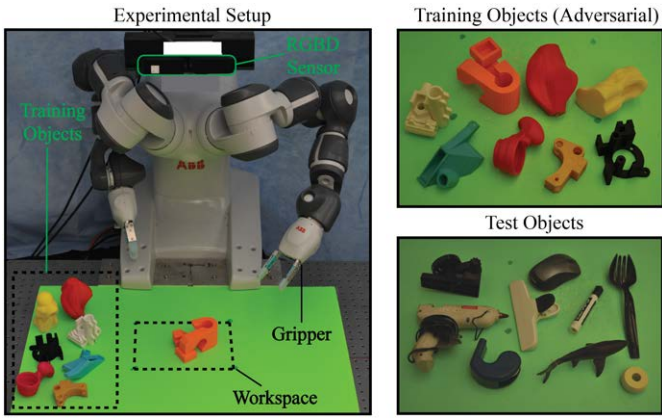


Fig. 5: (Left) The experimental platform for benchmarking grasping with the ABB YuMi. We registered the camera to the robot with a chessboard before each experiment. In each trial a human operator sampled an object pose by shaking the object in a box and placing it upside down in the workspace. We then took an RGB-D image with a Primsense Carmine 1.08, filled in the image using inpainting [25], segmented the object using color background subtraction, and formed a bounding box for the detected object. The grasp planner under evaluation then planned a gripper pose and the YuMi executed the grasp. Grasps were considered successful if the gripper held the object after lifting, transporting, and shaking the object. (Top-Right) The training set of 8 objects with adversarial geometric features such as smooth curved surfaces and narrow openings for grasping known objects. (Bottom-Right) The test set of 10 household objects not seen during training.

- 1) **Success Rate:** The percentage of grasps that were able to lift, transport, and hold a desired object after shaking.
- 2) **Precision:** The success rate on grasps that are have an estimated robustness higher than 50%. This measures performance when the robot can decide not to grasp an object, which could be useful when the robot has other actions (e.g. pushing) available.
- 3) **Robust Grasp Rate:** The percentage of planned grasps with an estimated robustness higher than 50%.
- 4) **Planning Time:** The time in seconds between receiving an image and returning a planned grasp.

B. Datasets

Fig. 5 illustrates the physical object datasets used in the benchmark:

- 1) **Train:** A validation set of 8 3D-printed objects with adversarial geometric features such as smooth, curved surfaces. This is used to set model parameters and to evaluate performance on known objects.
- 2) **Test:** A set of 10 household objects similar to models in Dex-Net 2.0 with various material, geometric, and specular properties. This is used to evaluate generalization to unknown objects.

We chose objects based on geometric features under three constraints: (a) small enough to fit within the workspace, (b) weight less than $0.25kg$, the payload of the YuMi, and (c) height from the table greater than $1.0cm$ due to a limitation of the silicone gripper fingertips.

We used four different GQ-CNN training datasets to study the effect on performance, each with a 80-20 image-wise training and validation split:




Fig. 6: Receiver operating characteristic comparing the performance of learning models on Adv-Synth. The GQ-CNN models all perform similarly and have a significantly higher true positive rate when compared to ML-RF and ML-SVM.

Model	Accuracy (%)
ML-SVM	89.7
ML-RF	90.5
GQ-S-Adv	97.8
GQ-L-Adv	97.8
GQ-Adv	98.1

TABLE II: The classification accuracy of each model on Adv-Synth. We see that the GQ-CNN methods have less than 2.5% test error while ML-RF and ML-SVM are closer to 10% error. Pretraining does not appear to affect performance.

- 1) **Adv-Synth:** Synthetic images and grasps for the adversarial objects in Train (189k datapoints).
- 2) **Adv-Phys:** Outcomes of executing random antipodal grasps with random gripper height and friction coefficient of $\mu = 0.5$ in 50 physical trials per object in Train (400 datapoints).
- 3) **Dex-Net-Small:** A subset of data from 150 models sampled uniformly from Dex-Net 2.0 (670k datapoints).
- 4) **Dex-Net-Large:** Data from all 1500 models in Dex-Net 2.0 (6.7m datapoints).

C. Grasp Planning Methods Used for Comparison

We compared a number of grasp planning methods on simulated and real data. We tuned the parameters of each method based on synthetic classification performance and physical performance on the training objects. All methods other than point cloud registration used the antipodal grasp sampling method described in Section V with the same set of parameters to generate candidate grasps, and each planner executes the highest-ranked grasp according to the method. Additional details on the methods and their parameters can be found in the supplemental file.

Image-based Grasp Quality Metrics (IGQ). We sampled a set of force closure grasp candidates by finding antipodal points on the object boundary [7] using edge detection and ranked grasps by the distance from the center of the jaws to the centroid of the object segmentation mask. We set the gripper depth using a fixed offset from the depth of the grasp center pixel.

Point-Cloud Registration (REG). We also compared with grasp planning based on point cloud registration, a state-of-the-art method for using precomputed grasps [14, 21]. We first coarsely estimated the object instance and pose based on the top 3 most similar synthetic images from Dex-Net 2.0, where similarity is measured as distance between AlexNet conv5 features [14, 35]. After coarse matching, we finetuned the pose of the object in the table plane using Iterated Closest Point [18, 28] with a point-to-plane cost. Finally, we retrieved the most robust gripper pose from Dex-Net 2.0 for the estimated object. The system had a median translational

	Comparisons of Methods						GQ-CNN Parameter Sensitivity					
	Random	IGQ	ML-RF	ML-SVM	REG	GQ-L-Adv	GQ-S-Adv	GQ-Adv	GQ-Adv-Phys	GQ-Adv-FC	GQ-Adv-LowU	GQ-Adv-HighU
Success Rate (%)	58±11	70±10	75±9	80±9	95±5	93±6	85±8	83±8	80±9	83±8	78±9	86±8
Precision (%)	N/A	N/A	100	100	N/A	94	90	91	80	89	90	92
Robust Grasp Rate (%)	N/A	N/A	5	0	N/A	43	60	44	100	89	53	64
Planning Time (sec)	N/A	1.9	0.8	0.9	2.6	0.8	0.9	0.8	0.8	0.7	0.8	0.9

TABLE III: Performance of grasp planning methods on the Train dataset with 95% confidence intervals for the success rate. Each method was tested for 80 trials (10 trials per object). Details on the methods used for comparison can be found in Section VI-C. We see that REG (point cloud registration) has the highest success rate at 95% but the GQ-L-Adv performs comparably at 93% and is $3\times$ faster. Performance of the GQ-CNN drops to 80% when trained on the Adv-Phys dataset (GQ-Adv-Phys), likely due to the small number of training examples, and drops to 78% when no noise is added to the images during training (GQ-Adv-LowU).

error of 4.5mm a median rotational error of 3.5° in the table plane for known objects.

Alternative Machine Learning Models (ML). We also compared the performance of a Random Forest with 200 trees of depth up to 10 (ML-RF) motivated by the results of [51] and a Support Vector Machine with the RBF kernel and a regularization parameter of 1 (ML-SVM) motivated by the results of [3, 50, 52]. For the RF we used the raw transformed images and gripper depths normalized by the mean and standard deviation across all pixels as features. For the SVM we used a Histogram of Oriented Gradients (HOG) [9] feature representation. Both methods were trained using scikit-learn on the Adv-Synth dataset.

Grasp Quality CNNs (GQ). We trained the GQ-CNN (abbrev. GQ) using the thresholded robust epsilon metric with $\delta = 0.002$ [26] for 5 epochs on Dex-Net-Large (all of Dex-Net 2.0) using Gaussian process image noise with standard deviation $\sigma = 0.005$. We used TensorFlow [1] with a batch size of 128, a momentum term of 0.9, and an exponentially decaying learning rate with step size 0.95. Training took approximately 48 hours on an NVIDIA GeForce 1080. The first layer of 7×7 convolution filters are shown in the right panel of Fig. 4, and suggest that the network learned finegrained vertical edge detectors and coarse oriented gradients. We hypothesize that vertical filters help to detect antipodal contact normals and the coarse oriented gradients estimate collisions.

To benchmark the architecture outside of our datasets, we trained on the Cornell Grasping Dataset [32] (containing 8,019 examples) and achieved a 93.0% recognition rate using grayscale images and an 80–20 imagewise training-validation split compared to 93.7% on RGB-D images in the original paper. We also trained several variants to evaluate sensitivity to several parameters:

Dataset Size. We trained a GQ-CNN on Dex-Net-Small for 15 epochs (GQ-S).

Amount of Pretraining We trained three GQ-CNNs on the synthetic dataset of adversarial training objects (Adv-Synth) to study the effect of pretraining with Dex-Net for a new, known set of objects. The model GQ-Adv was trained on only Adv-Synth for 25 epochs. The models GQ-L-Adv and GQ-S-Adv were initialized with the weights of GQ and GQ-S, respectively, and finetuned for 5 epochs on Adv-Synth.

Success Metric. We trained a GQ-CNN using probability of force closure thresholded at 25% (GQ-Adv-FC), which is a robust version of the antipodality metric of [16, 52], and

and labels for 400 random grasp attempts on the Train objects using a physical robot [33, 41] (GQ-Adv-Phys).

Noise Levels. We trained a GQ-CNN with zero noise $\sigma = 0$ (GQ-Adv-LowU) and high noise with $\sigma = 0.01$ (GQ-Adv-HighU).

D. Classification of Synthetic Data

The GQ-CNN trained on all of Dex-Net 2.0 had an accuracy of 85.7% on a held out validation set of approximately 1.3 million datapoints. Due to the memory and time requirements of training SVMs, we compared synthetic classification performance across methods on the smaller Adv-Synth dataset. Fig. 6 shows the receiver operating characteristic curve comparing the performance of GQ-L-Adv, GQ-S-Adv, GQ-Adv, ML-SVM, and ML-RF on a held-out validation set and Table II details the classification accuracy for the various methods. The GQ-CNNs outperformed ML-RF and ML-SVM, achieving near-perfect validation accuracy.

E. Performance Comparison on Known Objects

We evaluated the performance of the grasp planning methods on known objects from Train. Each grasp planner had 80 trials (10 per object). The left half of Table III compares the performance with other grasp planning methods and the right half compares the performance of the GQ-CNN variations. We found that GQ planned grasps $3\times$ faster than REG and achieved a high 93% success rate and 94% precision. The results also suggest that training on the full Dex-Net 2.0 dataset was necessary to achieve higher than 90% success.

F. Performance Comparison on Novel Objects

We also compared the performance of the methods on the ten novel test objects from Test to evaluate generalization to novel objects. Each method was run for 50 trials (5 per object). The parameters of each method were set based on Train object performance without knowledge of the performance on Test. Table IV details the results. GQ performed best with an 80% success rate and 100% precision (zero false positives over 29 grasps classified as robust).

G. Generalization Ability of the Dex-Net 2.0 Grasp Planner

We evaluated the generalization performance of GQ in 100 grasping trials on the 40 object test set illustrated in Fig. 7, which contains articulated (e.g. can opener) and deformable (e.g. washcloth) objects. We used the cross entropy method

	IGQ	REG	GQ-Adv-Phys	GQ-Adv	GQ-S	GQ
Success Rate (%)	60±13	52±14	68±13	74±12	72±12	80±11
Precision (%)	N/A	N/A	68	87	92	100
Robust Grasp Rate (%)	N/A	N/A	100	30	48	58
Planning Time (sec)	1.8	3.4	0.7	0.7	0.8	0.8

TABLE IV: Performance of grasp planning methods on our grasping benchmark with the test dataset of 10 household objects with 95% confidence intervals for the success rate. Each method was tested for 50 trials, and details on the methods used for comparison can be found in Section VI-C. GQ performs best in terms of success rate and precision, with 100% precision (zero false positives among 29 positive classifications). Performance decreases with smaller training datasets, but the GQ-CNN methods outperform the image-based grasp quality metrics (IGQ) and point cloud registration (REG).




Fig. 7: (Left) The test set of 40 household objects used for evaluating the generalization performance of the Dex-Net 2.0 grasp planner. The dataset contains rigid, articulated, and deformable objects. (Right) The experimental setup for order fulfillment with the ABB YuMi. The goal is to grasp and transport three target objects to a shipping container (box on right).

(CEM) [33], which iteratively samples a set of candidate grasps and re-fits the candidate grasp distribution to the grasps with the highest predicted robustness, in order to find better maxima of the robust grasping policy. More details can be found in the supplemental file. The CEM-augmented Dex-Net 2.0 grasp planner achieved 94% success and 99% precision (68 successes out of 69 grasps classified as robust), and it took an average of 2.5s to plan grasps.

H. Application: Order Fulfillment

To demonstrate the modularity of the Dex-Net 2.0 grasp planner, we used it in an order fulfillment application with the ABB YuMi. The goal was to grasp and transport a set of three target objects to a shipping box in the presence of three distractor objects when starting with the objects in a pile on a planar worksurface, illustrated in Fig. 7. Since the Dex-Net 2.0 grasp planner assumes singulated objects, the YuMi first separated the objects using a policy learned from human demonstrations mapping binary images to push locations [31]. When the robot detected an object with sufficient clearance from the pile, it identified the object based on color and used GQ-L-Adv to plan a robust grasp. The robot then transported the object to either the shipping box or a reject box, depending on whether or not the object was a distractor. The system successfully placed the correct objects in the box on 4 out of 5 attempts and was successful in grasping on 93% of 27 total attempts.

I. Failure Modes

Fig. 8 displays some common failures of the GQ-CNN grasp planner. One failure mode occurred when the RGB-D




Fig. 8: Four examples of failed grasps planned using the GQ-CNN from Dex-Net 2.0. The most common failure modes were related to: (left) missing sensor data for an important part of the object geometry, such as thin parts of the object surface, and (right) collisions with the object that are misclassified as robust.

sensor failed to measure thin parts of the object geometry, making these regions seem accessible. A second type of failure occurred due to collisions with the object. It appears that the network was not able to fully distinguish collision-free grasps in narrow parts of the object geometry. This suggests that performance could be improved with more accurate depth sensing and using analytic methods to prune grasps in collision.

VII. DISCUSSION AND FUTURE WORK

We developed a Grasp Quality Convolutional Neural Network (GQ-CNN) architecture that predicts grasp robustness from a point cloud and trained it on Dex-Net 2.0, a dataset containing 6.7 million point clouds, parallel-jaw grasps, and robust grasp metrics. In over 1,000 physical evaluations, we found that the Dex-Net 2.0 grasp planner is as reliable and 3× faster a method based on point cloud registration, and had 99% precision on a test set of 40 novel objects.

In future work, our goal is to approach 100% success on known objects by using active learning to adaptively acquire grasps using a policy initialized with a GQ-CNN. Additionally, we plan to extend the method to grasp objects in clutter [16, 33] by using simulated piles of rigid objects from Dex-Net and by augmenting the grasping policy with an option to push and separate objects when no robust grasp is available. We also intend to extend the method to use point clouds from multiple viewpoints and in grasping tasks with sequential structure, such as regrasping for assembly. Furthermore, we plan to release a subset of our code, dataset, and the trained GQ-CNN weights to facilitate further research and comparisons.

ACKNOWLEDGMENTS

This research was performed at the AUTOLAB at UC Berkeley in affiliation with the Berkeley AI Research (BAIR) Lab, the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS People and Robots (CPAR) Initiative. The authors were supported in part by the U.S. National Science Foundation under NRI Award IIS-1227536: Multilateral Manipulation by Human-Robot Collaborative Systems, the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program, the Berkeley Deep Drive (BDD) Program, and by donations from Siemens, Google, Cisco, Autodesk, IBM, Amazon Robotics, and Toyota Robotics Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Sponsors. We thank our colleagues who provided helpful feedback, code, and suggestions, in particular Pieter Abbeel, Ruzena Bajcsy, Brenton Chu, Roy Fox, David Gealy, Ed Johns, Sanjay Krishnan, Animesh Garg, Sergey Levine, Pusong Li, Matt Matl, Stephen McKinley, Andrew Reardon, Vishal Satish, Sammy Staszak, and Nan Tian.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhiheng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [2] Ravi Balasubramanian, Ling Xu, Peter D Brook, Joshua R Smith, and Yoky Matsuoka. Physical human interactive guidance: Identifying grasping principles from human-planned grasps. *IEEE Trans. Robotics*, 28(4):899–910, 2012.
- [3] Jeannette Bohg and Danica Kragic. Learning grasping points with shape context. *Robotics and Autonomous Systems*, 58(4):362–377, 2010.
- [4] Jeannette Bohg, Antonio Morales, Tamim Asfour, and Danica Kragic. Data-driven grasp synthesis survey. *IEEE Trans. Robotics*, 30(2):289–309, 2014.
- [5] Peter Brook, Matei Ciocarlie, and Kaijen Hsiao. Collaborative grasp planning with multiple object representations. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 2851–2858. IEEE, 2011.
- [6] Lillian Chang, Joshua R Smith, and Dieter Fox. Interactive singulation of objects from a pile. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 3875–3882. IEEE, 2012.
- [7] I-Ming Chen and Joel W Burdick. Finding antipodal point grasps on irregularly shaped objects. *IEEE Trans. Robotics and Automation*, 9(4):507–512, 1993.
- [8] Matei Ciocarlie, Kaijen Hsiao, Edward Gil Jones, Sachin Chitta, Radu Bogdan Rusu, and Ioan A Şucan. Towards reliable grasping and manipulation in household environments. In *Experimental Robotics*, pages 241–252. Springer, 2014.
- [9] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 886–893. IEEE, 2005.
- [10] Renaud Detry, Carl Henrik Ek, Marianna Madry, and Danica Kragic. Learning a dictionary of prototypical grasp-predicting parts from grasping experience. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 601–608. IEEE, 2013.
- [11] Andreas Eitel, Jost Tobias Springenberg, Luciano Spinello, Martin Riedmiller, and Wolfram Burgard. Multimodal deep learning for robust rgb-d object recognition. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 681–687. IEEE, 2015.
- [12] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.
- [13] Ken Goldberg, Brian V Mirtich, Yan Zhuang, John Craig, Brian R Carlisle, and John Canny. Part pose statistics: Estimators and experiments. *IEEE Trans. Robotics and Automation*, 15(5):849–857, 1999.
- [14] Corey Goldfeder and Peter K Allen. Data-driven grasping. *Autonomous Robots*, 31(1):1–20, 2011.
- [15] Corey Goldfeder, Matei Ciocarlie, Hao Dang, and Peter K Allen. The columbia grasp database. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 1710–1716. IEEE, 2009.
- [16] Marcus Gualtieri, Andreas ten Pas, Kate Saenko, and Robert Platt. High precision grasp pose detection in dense clutter. *arXiv preprint arXiv:1603.01564*, 2016.
- [17] Menglong Guo, David V Gealy, Jacky Liang, Jeffrey Mahler, Aimee Goncalves, Stephen McKinley, and Ken Goldberg. Design of parallel-jaw gripper tip surfaces for robust grasping. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017.
- [18] Saurabh Gupta, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Aligning 3d models to rgb-d images of cluttered scenes. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4731–4740, 2015.
- [19] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pages 1026–1034, 2015.
- [21] Carlos Hernandez, Mukunda Bharatheesha, Wilson Ko, Hans Gaiser, Jethro Tan, Kanter van Deurzen, Maarten de Vries, Bas Van Mil, Jeff van Egmond, Ruben Burger, et al. Team delft’s robot winner of the amazon picking challenge 2016. *arXiv preprint arXiv:1610.05514*, 2016.
- [22] Alexander Herzog, Peter Pastor, Mrinal Kalakrishnan, Ludovic Righetti, Jeannette Bohg, Tamim Asfour, and Stefan Schaal. Learning of grasp selection based on shape-templates. *Autonomous Robots*, 36(1-2):51–65, 2014.
- [23] Stefan Hinterstoisser, Stefan Holzer, Cedric Cagniart, Slobodan Ilic, Kurt Konolige, Nassir Navab, and Vincent Lepetit. Multi-modal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, pages 858–865. IEEE, 2011.
- [24] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Proc. Advances in Neural Information Processing Systems*, pages 2017–2025, 2015.
- [25] Edward Johns, Stefan Leutenegger, and Andrew J Davison. Deep learning a grasp function for grasping under gripper pose uncertainty. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 4461–4468. IEEE, 2016.
- [26] Daniel Kappler, Jeannette Bohg, and Stefan Schaal. Leveraging big data for grasp planning. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2015.
- [27] Alexander Kasper, Zhixing Xue, and Rüdiger Dillmann. The kit object models database: An object model database for object recognition, localization and manipulation in service robotics. *Int. Journal of Robotics Research (IJRR)*, 31(8):927–934, 2012.
- [28] Ben Kehoe, Akihiro Matsukawa, Sal Candido, James Kuffner, and Ken Goldberg. Cloud-based robot grasping with the google object recognition engine. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 4263–4270. IEEE, 2013.
- [29] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [30] Michael Laskey, Jeffrey Mahler, Zoe McCarthy, Florian T Pokorny, Sachin Patil, Jur van den Berg, Danica Kragic, Pieter Abbeel, and Ken Goldberg. Multi-armed bandit models for 2d grasp planning with uncertainty. In *Proc. IEEE Conf. on Automation Science and Engineering (CASE)*. IEEE, 2015.

- [31] Michael Laskey, Caleb Chuck, Jonathan Lee, Jeffrey Mahler, Sanjay Krishnan, Kevin Jamieson, Anca Dragan, and Ken Goldberg. Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2017.
- [32] Ian Lenz, Honglak Lee, and Ashutosh Saxena. Deep learning for detecting robotic grasps. *Int. Journal of Robotics Research (IJRR)*, 34(4-5):705–724, 2015.
- [33] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *arXiv preprint arXiv:1603.02199*, 2016.
- [34] Jeffrey Mahler, Brian Hou, Sherdil Niyaz, Florian T Pokorny, Ramu Chandra, and Ken Goldberg. Privacy-preserving grasp planning in the cloud. In *Proc. IEEE Conf. on Automation Science and Engineering (CASE)*, pages 468–475. IEEE, 2016.
- [35] Jeffrey Mahler, Florian T Pokorny, Brian Hou, Melrose Roderick, Michael Laskey, Mathieu Aubry, Kai Kohlhoff, Torsten Kröger, James Kuffner, and Ken Goldberg. Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2016.
- [36] Tanwi Mallick, Partha Pratim Das, and Arun Kumar Majumdar. Characterizations of noise in kinect depth images: A review. *IEEE Sensors Journal*, 14(6):1731–1740, 2014.
- [37] John W Miller, Rod Goodman, and Padhraic Smyth. On loss functions which minimize to conditional expected values and posterior probabilities. *IEEE Transactions on Information Theory*, 39(4):1404–1408, 1993.
- [38] Luis Montesano and Manuel Lopes. Active learning of visual descriptors for grasping using non-parametric smoothed beta distributions. *Robotics and Autonomous Systems*, 60(3):452–462, 2012.
- [39] Anh Nguyen, Dimitrios Kanoulas, Darwin G Caldwell, and Nikos G Tsagarakis. Detecting object affordances with convolutional neural networks. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 2765–2770. IEEE, 2016.
- [40] John Oberlin and Stefanie Tellex. Autonomously acquiring instance-based object models from experience. In *Int. S. Robotics Research (ISRR)*, 2015.
- [41] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2016.
- [42] Lerrel Pinto, James Davidson, and Abhinav Gupta. Supervision via competition: Robot adversaries for learning tasks. *arXiv preprint arXiv:1610.01685*, 2016.
- [43] Florian T Pokorny and Danica Kragic. Classical grasp quality evaluation: New algorithms and theory. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 3493–3500. IEEE, 2013.
- [44] Lorenzo Porzi, Samuel Rota Buló, Adrian Penate-Sanchez, Elisa Ricci, and Francesc Moreno-Noguer. Learning depth-aware deep representations for robotic perception. *IEEE Robotics & Automation Letters*, 2016.
- [45] Domenico Prattichizzo and Jeffrey C Trinkle. Grasping. In *Springer handbook of robotics*, pages 671–700. Springer, 2008.
- [46] Joseph Redmon and Anelia Angelova. Real-time grasp detection using convolutional neural networks. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 1316–1322. IEEE, 2015.
- [47] Alberto Rodriguez, Matthew T Mason, and Steve Ferry. From caging to grasping. *Int. Journal of Robotics Research (IJRR)*, page 0278364912442972, 2012.
- [48] Fereshteh Sadeghi and Sergey Levine. (cad)2 rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.
- [49] Renato F Salas-Moreno, Richard A Newcombe, Hauke Strasdat, Paul HJ Kelly, and Andrew J Davison. Slam++: Simultaneous localisation and mapping at the level of objects. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1352–1359, 2013.
- [50] Ashutosh Saxena, Justin Driemeyer, and Andrew Y Ng. Robotic grasping of novel objects using vision. *The International Journal of Robotics Research*, 27(2):157–173, 2008.
- [51] Daniel Seita, Florian T Pokorny, Jeffrey Mahler, Danica Kragic, Michael Franklin, John Canny, and Ken Goldberg. Large-scale supervised learning of the grasp robustness of surface patch pairs. In *Proc. IEEE Int. Conf. on Simulation, Modeling, and Programming of Autonomous Robots (SIMPAR)*. IEEE, 2016.
- [52] Andreas ten Pas and Robert Platt. Using geometry to detect grasp poses in 3d point clouds. In *Intl Symp. on Robotics Research*, 2015.
- [53] Eric Tzeng, Coline Devin, Judy Hoffman, Chelsea Finn, Pieter Abbeel, Sergey Levine, Kate Saenko, and Trevor Darrell. Adapting deep visuomotor representations with weak pairwise constraints. In *Workshop on the Algorithmic Foundation of Robotics (WAFR)*, 2016.
- [54] Jacob Varley, Jonathan Weisz, Jared Weiss, and Peter Allen. Generating multi-fingered robotic grasps via deep learning. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 4415–4420. IEEE, 2015.
- [55] Matthew Veres, Medhat Moussa, and Graham W Taylor. Modeling grasp motor imagery through deep conditional generative models. *arXiv preprint arXiv:1701.03041*, 2017.
- [56] Jonathan Weisz and Peter K Allen. Pose error robust grasping from contact wrench space metrics. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 557–562. IEEE, 2012.
- [57] Walter Wohlkinger, Aitor Aldoma, Radu B Rusu, and Markus Vincze. 3dnet: Large-scale object class recognition from cad models. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, pages 5384–5391. IEEE, 2012.
- [58] Ziang Xie, Arjun Singh, Justin Uang, Karthik S Narayan, and Pieter Abbeel. Multimodal blending for high-accuracy instance recognition. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 2214–2221. IEEE, 2013.
- [59] Andy Zeng, Kuan-Ting Yu, Shuran Song, Daniel Suo, Ed Walker Jr, Alberto Rodriguez, and Jianxiang Xiao. Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge. *arXiv preprint arXiv:1609.09475*, 2016.
- [60] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. *arXiv preprint arXiv:1609.05143*, 2016.