Science Robotics
AAAS

robotics.sciencemag.org/cgi/content/full/4/26/eaau4984/DC1

Supplementary Materials for

Learning ambidextrous robot grasping policies

Jeffrey Mahler*, Matthew Matl, Vishal Satish, Michael Danielczuk, Bill DeRose, Stephen McKinley, Ken Goldberg

*Corresponding author. Email: jmahler@berkeley.edu

Published 16 January 2019, *Sci. Robot.* **4**, eaau4984 (2019)
DOI: 10.1126/scirobotics.aau4984

The PDF file includes:

Text
Fig. S1. Analysis of features learned by the GQ-CNNs from the ambidextrous grasping policy.
Fig. S2. Per-object reliability of each policy on each test object.
Fig. S3. Difficulty of each object from the test object datasets characterized by the overall reliability averaged across methods.
Legend for movie S1
Legend for raw videos from experiments

Other Supplementary Material for this manuscript includes the following:

(available at robotics.sciencemag.org/cgi/content/full/4/26/eaau4984/DC1)

Movie S1 (.mp4 format). Summary.
Raw data, code for data analysis, videos, and listing of objects used in experiments (.zip file)

Text

For raw data and detailed results, please see http://bit.ly/2M3Dx1e.

Characterization of Grasping Benchmark

We characterize the Universal Picking benchmark in terms of several important variables common to robot grasping and manipulation (*58*).

Procedure. The test procedure for each policy consists of five independent trials in which the bin is filled with a random configuration of objects and the robot attempts to pick each object from the bin and transport it to a receptacle. In each trial, the operator sets a dataset of objects in the bin by shaking the objects in a box and placing the box upside down in the bin. On each grasp attempt, the robot policy receives as input a point cloud of the objects in the bin and returns a grasp action for one of the grippers. Then the robot moves to the target pose, attempts to grasp the object, and attempts to move the object to a receptacle on the side of the bin. If the robot successfully lifts and transport one object to the receptacle, then the operator labels the trial a success. A trial is considered complete after either all objects are removed, 75 total attempts, or 10 consecutive failures.

Object Range. Experiments use a dataset of 75 objects chosen to reflect a diverse range of shapes, sizes, and material properties. The dataset is broken into three difficulty levels. Difficulty level 1 contains prismatic and circular solids. Difficulty level 2 contains common household objects with complex geometry including toys, tools, and "blisterpack" objects. Difficulty level 3 contains objects with adversarial geometry and material properties such as detailed 3D-printed industrial parts and deformable objects. A complete listing of objects used in experiments, with purchase links, is available at http://bit.ly/2wGPsvZ.

Success Metrics. A grasp is considered successful if it lifts and transports exactly one object from the heap to a receptacle.

Computer System. All experiments run on a Desktop running Ubuntu 16.04 with a 3.4 GHz Intel Core i7-6700 Quad-Core CPU and an NVIDIA Titan Xp GPU.

Robot Arm. The benchmark uses an ABB YuMi industrial robot with two 7-DOF arms. YuMi has a repeatability of $0.02mm$, a reach of $559mm$, and a maximum tool speed of $1.5m/s$.

Robot Grippers / Hands. The robot uses a custom vacuum suction cup gripper and a parallel-jaw gripper with custom silicone fingertips (*57*). The suction gripper consists of a $20mm$ diameter silicone single-bellow suction cup seated in a 3D printed housing mounted to the end of the right arm. The vacuum is created by supplying compressed air from a Jun Air 18-40 quiet air compressor to a VacMotion MSV-35 vacuum generator. The payload of the suction system was approximately $0.9kg$ with a vacuum flow of approximately 8 standard cubic feet per minute.

Sensors. The robot plans grasps based on 3D point clouds from a single Photoneo PhoXi S industrial depth camera. The depth camera is mounted in a fixed location above the bin. The bin is mounted on a set of LoadStar load cells that measure the weight of objects in the bin with a resolution of approximately $5g$.

Controller. The robot moves the arm to grasping positions using the ABB RAPID motion planner and controller for position control of the tool with linear movements.

Lighting. The workspace is illuminated by several overhead interior lights and ambient sunlight from a single window to the outdoors.

Application. The application is Universal Picking, in which a robot rapidly grasps and transports objects from a bin to a receptacle.

Detailed Parameters of Dataset Generation

The YAML configuration files listing all parameters of dataset generation are included in the supplement. We detail the most important parameters here.

State Distribution. The initial state distribution $\xi(\mathbf{x}_0)$ is the product of distributions on (*26*):

1. *Object Count (m):* Truncated Poisson distribution with mean $\lambda = 5$ with a minimum of 3 objects and a maximum of 10 objects (to prevent excessive runtime).

2. *Object Heap ($\mathcal{O}$):* Uniform distribution over $1,664$ 3D object models from Thingiverse and the pose from which each model is dropped into a heap. Objects are sampled without replacement. The object drop pose is sampled by selecting an orientation uniformly at random and a position uniformly from the 2D plane $[-0.125m, -0.125m] \times [0.125m, 0.125m]$ at height $z = 0.15m$ above the a 3D bin. The drop is simulated using pybullet (*51*): until either (a) the positional velocity drops below a threshold $0.005m/s$ and the angular velocity drops below a threshold $0.1rad/sec$ or (b) 500 timesteps have elapsed (to prevent excessive runtime).

3. *Depth Camera ($\mathcal{C}$):* The camera pose is sampled uniformly from a set of spherical coordinates: $r \in [0.7m, 0.9m], \theta \in [250°, 290°], \varphi \in [0°, 10°]$. The center of the sphere is selected randomly from the planar base of the bin: $x \in [-0.05m, 0.05m], y \in [-0.05m, 0.05m]$. The intrinsic parameters are sampled from uniform distributions on the following sets: focal length $f \in [550px, 600px]$, camera optical center pixel $c_x \in [257.5px, 262.5px], c_y \in [257.5px, 262.5px]$. Images are rendered at half resolution with dimensions 500×500.

4. *Coulomb Friction (γ):* Truncated Gaussian constrained to $[0, 1]$ with a mean of 0.8

(based on the coefficient of friction between silicone and plastic) and a variance of 0.01.

Reward Distribution. Binary rewards occur when a quasi-static equilibrium is feasible between the grasp and an external wrench perturbation (e.g. due to gravity or inertia). Let $\mathcal{O}_i \in \mathbf{x}_t$ be an object contacted by the gripper when executing action $\mathbf{u}_t$. Then we measure grasp success with a binary-valued metric $S(\mathbf{x}_t, \mathbf{u}_t) \in \{0, 1\}$ that measures whether or not:

- The gripper geometry in the pose specified by $\mathbf{u}_t$ is collision-free from all objects and obstacles. We use a known 3D CAD model for the gripper.
- The gripper contacts exactly one object $\mathcal{O}_i$ when executing the grasp parameterized by $\mathbf{u}_t$.
- The grasp can resist a random disturbing force and torque (wrench) $\mathbf{w}_t = \mathbf{w}_g + \epsilon_w$ on the grasped object with over 50% probability, where $\mathbf{w}_g$ is the fixed wrench due to gravity and ϵ_w is a random wrench sampled from a zero-mean Gaussian $\mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I})$. We set $\sigma_w = 0.01$.

Given an object consisting of a geometry $\mathcal{M}$ in pose $\mathbf{T}_o$, the gripper g (geometry and physical parameters such as friction) and grasp pose $\mathbf{T}_g$ are used to determine the contacts $\mathbf{c}$, or set of points and normals between the fingers and object. This set of contacts is used to compute the set of wrenches Λ that the grasp can apply to the object under quasi-static physics. We use a soft finger contact model (*1*) for the parallel-jaw gripper and a compliant suction ring contact model (*27*) for the suction cup gripper. The grasp reward $R = 1$ if the probability of wrench resistance is greater than a threshold over $M = 10$ samples from the stochastic model. During sampling, we randomize over the object mass and gripper pose. The object mass is sampled from a Gamma distribution with mean $0.25kg$ and variance $0.05kg^2$. The parallel-jaw gripper pose is sampled from a zero-mean Gaussian with a positional standard deviation of $2.5mm$. The suction gripper pose is sampled from a zero-mean Gaussian with a positional standard deviation of $0.1mm$.

Data Collection Policy. The dataset collection policy $\tau(\mathbf{u}_t \mid \mathbf{x}_t, \mathbf{y}_t)$ samples a mixture of actions from the point cloud and from an algorithmic supervisor $\Omega(\mathbf{x})$ that guides data toward successful grasps:

$$\tau(\mathbf{u}_t \mid \mathbf{x}_t, \mathbf{y}_t) = \begin{cases} \Omega(\mathbf{x}_t) & \text{with prob. } \epsilon \\ Unif(\mathcal{U}_g(\mathbf{y}_t)) & \text{otherwise} \end{cases}$$

We use $\epsilon = 1\%$ to favor actions sampled from the policy's own action space, which may reduce covariate shift.

The set $\mathcal{U}_g(\mathbf{y})$ is the set of candidate actions sampled from the point cloud with equal numbers of suction and parallel-jaw grasp actions. For the parallel-jaw gripper, grasps are sampled using the antipodal grasp sampling method of Dex-Net 2.0 (*19*) with the same parameters. For the suction cup gripper, grasps are sampled uniformly from the set of points above the bottom of the bin with the same parameters as Dex-Net 3.0 (*27*).

The algorithmic supervisor $\Omega(\mathbf{x}_t)$ computes robust quality for a set of grasps on each 3D object model, modeling uncertainty in gripper pose, friction, mass, and the direction of gravity at runtime. Quality is computed with robust wrench resistance using $M = 25$. The supervisor samples the gripper pose from a Gaussian random variable. The parallel-jaw gripper pose is sampled from a zero-mean Gaussian with a positional standard deviation of $2.5mm$. The suction gripper pose is sampled from a zero-mean Gaussian with a positional standard deviation of $0.1mm$. Friction is modeled as a Truncated Gaussian constrained to $[0, 1]$ with a mean of 0.8 (based on the coefficient of friction between silicone and plastic) and a variance of 0.01. The object mass is sampled from a Gamma distribution with mean

$0.25kg$ and variance $0.05kg^2$. The direction of gravity is chosen based on the gripper. For the parallel-jaw gripper, the direction of gravity is sampled uniformly at random from spherical coordinates over the unit 2D sphere. For the suction cup gripper, the direction of gravity is sampled within a cone of angle $\pi/6$ from the suction approach axis using rejection sampling. The supervisor sampled grasps uniformly at random from the set with quality greater than a threshold of 75%.

Parameters of Grasping Policy

The policy optimizes for the highest quality grasp for each gripper separately using the Cross Entropy Method (CEM) (*10, 19, 27, 56*), and then selects the grasp with the highest estimated quality across the grippers. The parameters of the CEM-based grasping policy include: a number of iterations $m = 3$, an initial candidate set size of $n = 250$, a number of resamples $c = 50$, a number of Gaussian Mixture Model components $k = 5$, and an elite-set-percentile $\gamma = 0.25$ using an identical implementation to Dex-Net 2.0 (*19*) and Dex-Net 3.0 (*27*). The parameters were the same for both grippers.

Figures

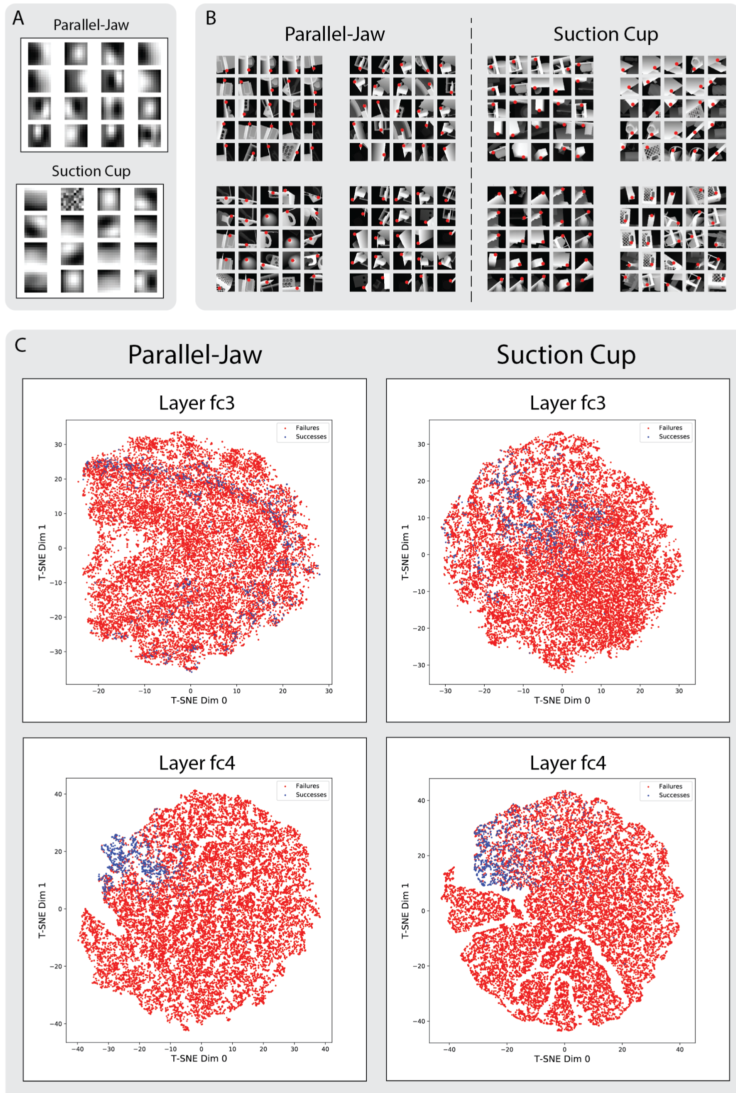

Fig. S1. Analysis of features learned by the GQ-CNNs from the ambidextrous grasping policy. (**A**) First layer of 16 9×9 convolutional filters from each network. The filters appear to respond to oriented gradients. The suction network contains one random filter, which may be a "dead neuron," or neuron that produces zero output on all grasps the training dataset. (**B**) Images from the Dex-Net 4.0 synthetic validation set that trigger the maximum responses for neurons from the highest convolutional layer. A subset of four neurons are shown for each network. The parallel-jaw neurons appear to respond to thin, oriented bars which may afford robust grasps, although some neurons respond to corners and are less interpretable. The suction neurons appear to respond to sloped flat surfaces and oriented object edges and corners, which may be useful for estimating object centroids. (**C**) t-Stochastic Neighbor Embedding (t-SNE) plot for a randomly selected subset of images from the Dex-Net 4.0 synthetic validation set. Each point corresponds to a candidate grasp and point cloud observation. In both networks, the positive examples (blue) and negative examples (red) are not well-separated until the last fully connected layer, fc4. Furthermore, the datapoints do not show significant clustering. The plot was generated using Barnes Hut t-SNE with a perplexity of 25.

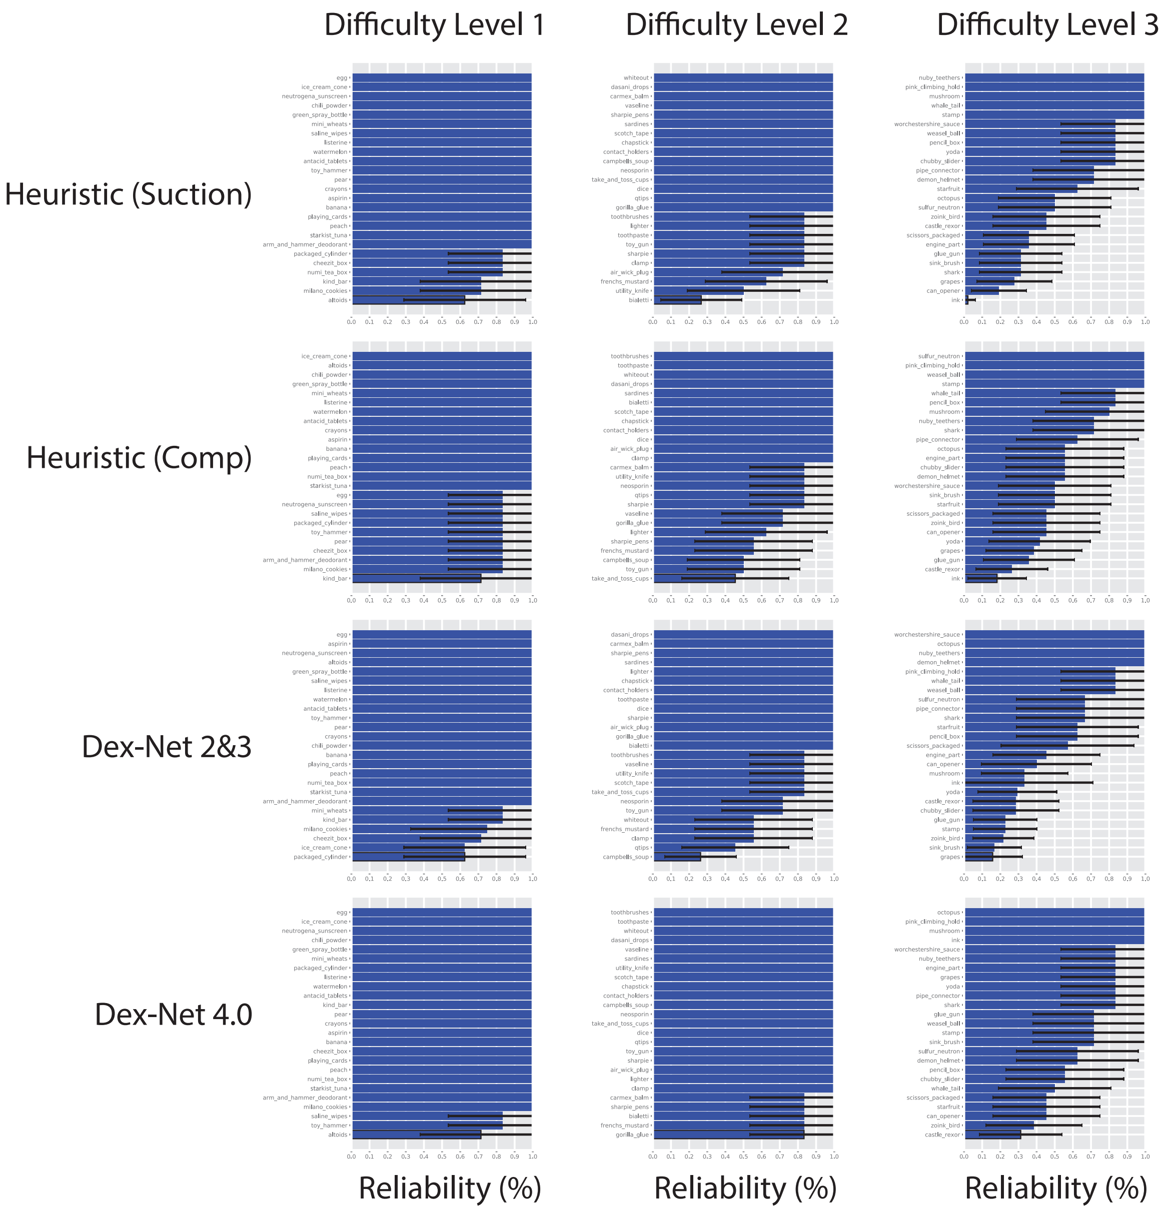

Fig. S2. Per-object reliability of each policy on each test object. For each policy, objects are listed from top to bottom in order of decreasing reliability to facilitate characterization of the most difficult objects. Error bars show the 95% confidence interval on reliability using the standard error of the mean. The differences in reliability between policies on the Level 1 objects cannot be attributed to any single challenging objects. There are several challenging Level 2 objects for each baseline: the bialetti espresso blisterpack (Suction Heuristic), the take-and-toss cups toy (Composite Heuristic), and the campbells soup can (DexNet 2&3). The Level 3 objects are challenging for each policy, with the baselines having difficulty on the ink, glue gun, and grapes.

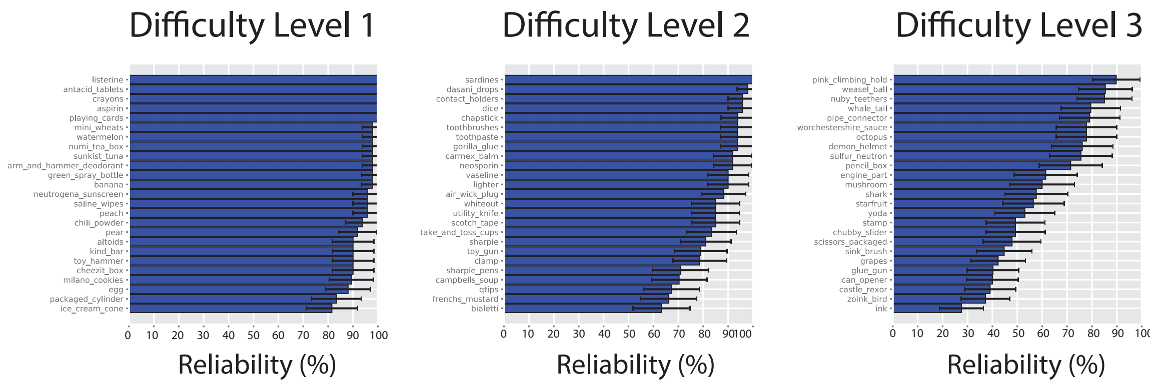

Fig. S3. Difficulty of each object from the test object datasets characterized by the overall reliability averaged across methods. Error bars display the standard error of the mean on the reliability for each object averaged across all policies benchmarked in the experiments. The Level 1 objects all have above 80% reliability, with the most difficult object being the toy ice cream cone. The Level 2 objects contain several objects with reliability between 60% and 80%. These may be categorized as heavy objects (e.g., q-tips, mustard), objects with "blisterpack" and transparent plastic (e.g., bialetti, sharpie pens), and objects with complex geometry (e.g., toy gun, clamp). The Level 3 objects have considerably lower reliability. The most difficult objects were: the bottle of ink, which was difficult to grasp in corners of the bin, the "zoink" bird dog toy, which had significant areas of porous material that could not be suctioned, and castle rexor, a 3D printed object with very finely detailed geometry.

Movie S1. Summary.

The movie summarizes the motivation, methods, and results of the paper and is best viewed in a small window due to low resolution.

The policy plans the highest quality grasp for each gripper using iterative optimization. This animation depicts the set of grasps sampled by the Cross Entropy Method (CEM) for derivative-free optimization as the policy searches for the grasp with highest predicted quality for each gripper. Grasps are colored based on quality (red: $Q = 0\%$, blue: $Q = 100\%$). As the number of iterations increases, the sample set converges to a set of high quality grasps. The final grasp is selected by comparing the quality of the best grasp planned for each gripper.

Example experimental trial with novel test objects from Level 2. This shows a video for grasp attempts on an ABB YuMi robot that were used to report the results for the Dex-Net 4.0 policy on the Level 2 objects.

The left hand side shows a point cloud observation rendered as a depth image. On each grasp attempt, the planned grasp is projected into the image and rendered at the corresponding pixel. Dots corresponding to suction grasps and a lines correspond to a parallel-jaw grasps. In addition, the trial number, number of grasp attempts, number of objects picked, and reward for each attempt are plotted.

The right hand side shows a color video of the grasp being executed on the robot, taken with an overhead webcam.

The overall reliability of Dex-Net 4.0 on the Level 2 objects was 95% with 125 successful grasps in 131 attempts.

Raw videos from experiments

Footage from each trial used to report results in the paper was automatically compiled into a set of videos. The videos can be found at http://bit.ly/2M3Dx1e in zip file *aau4894_supplementary_material.zip*. In the unzipped file, videos are in the folder *results/videos*.

Each video shows all grasp attempts made by a single candidate policy for the five trials of bin picking described in the paper. The video filenames contain the name of the policy that appears in the footage (e.g., “Dex-Net 4.0.mp4”).

These are exactly the grasp attempts used to derive the results presented. Please keep in mind that the data is imperfect and subject to some errors due to human labeling, but we have made our best effort to remove these cases.

In the video, the left hand side shows a point cloud observation rendered as a depth image. On each grasp attempt, the planned grasp is projected into the image and rendered with a dot corresponding to a suction grasp and a line corresponding to a parallel-jaw grasp. In addition, the trial number, number of grasp attempts, number of objects picked, and reward for each attempt are plotted.

When the robot stops moving for a long period of time, a red "COLLISION" is rendered on the screen to indicate one of two collision failures: a joint collision (planning failure of the ABB RAPID system) or a collision with the environment that jams the robot.

The right hand side shows a color video of the grasp being executed on the robot, taken with an overhead web cam.