

# Multi-Fingered Active Grasp Learning

Qingkai Lu<sup>1</sup>, Mark Van der Merwe<sup>1</sup>, and Tucker Hermans<sup>1,2</sup>

**Abstract**—Learning-based approaches to grasp planning are preferred over analytical methods due to their ability to better generalize to new, partially observed objects. However, data collection remains one of the biggest bottlenecks for grasp learning methods, particularly for multi-fingered hands. The relatively high dimensional configuration space of the hands coupled with the diversity of objects common in daily life requires a significant number of samples to produce robust and confident grasp success classifiers. In this paper, we present the first active deep learning approach to grasping that searches over the grasp configuration space and classifier confidence in a unified manner. We base our approach on recent success in planning multi-fingered grasps as probabilistic inference with a learned neural network likelihood function. We embed this within a multi-armed bandit formulation of sample selection. We show that our active grasp learning approach uses fewer training samples to produce grasp success rates comparable with the passive supervised learning method trained with grasping data generated by an analytical planner. We additionally show that grasps generated by the active learner have greater qualitative and quantitative diversity in shape.

## I. INTRODUCTION

Learning-based grasp planning [1–9] has become popular over the past decade, because of its ability to generalize well to novel objects with only partial-view object information [10]. These approaches require large amounts of data for training, particularly those that utilize deep neural networks [8, 9, 11–13]. However, large scale data collection remains a challenge for multi-fingered grasping, because (1) objects common in daily life exhibit large variation in terms of geometry, texture, inertial properties, and appearance; and (2) the relatively high dimension of multi-fingered grasp configurations, (e.g. 22 dimensions for the configuration of hand and wrist pose in this paper).

Random sampling is a common approach for data collection of two-finger gripper grasping [5, 7, 14]. Considering the sparsity of successful grasps in the high dimensional configuration space of multi-fingered grasping, it is not practical to collect enough successful grasps using random sampling alone. Eppner et al. [15] exhaustively sample more than 1 billion grasps for each of 21 objects from the YCB data set to improve the understanding of data generation for 6-DOF, parallel jaw grasp learning algorithms. However, this approach would not scale well to multi-fingered grasping due to its high dimensionality.

Existing deep learning work [8, 11, 12] for multi-fingered grasping collect training grasps using analytical or heuristic grasp planners to combat the issues of sampling in high-dimensional spaces. While these passive supervised grasp learning techniques generalize well over the space covered by the training data, the grasp planners used for training bias data collection to only cover a subspace of all feasible grasps configurations. Additionally grasps generated for different object geometries tend to be quite similar in shape.

We propose an active learning approach to interactively learn a grasp model that better covers the grasp configuration space across different objects using fewer samples compared with a passive, supervised grasp learner. Instead of passively inducing a hypothesis to explain the available training data as in standard supervised learning, active learning develops and tests new hypotheses continuously and interactively. Active learning is most appropriate when 1) unlabeled data samples are numerous, 2) a lot of labeled data are needed to train an accurate supervised learning system, and 3) data samples can be easily collected or synthesized [16]. Grasp learning satisfies each of these conditions: 1) there are infinitely many possible grasps, 2) a large number of labeled training samples are necessary to cover the space [15], and 3) the robot is its own oracle—it can try a grasps and automatically detecting success or failure without human labeling.

We propose modeling active grasp learning as a multi-armed bandit problem, which is designed to improve the grasp success classifier and enable the grasp model to cover the space of grasp configurations and objects as much as possible. Our proposed method is fundamentally different from existing bandit-based grasping work [17–19], which treat each possible grasp as one arm and perform grasp planning separately for different objects. We instead use the bandit-framework to select between three qualitatively different exploration strategies each implemented using the grasp planning as probabilistic inference formulation [13].

We perform grasp planning using our actively learned grasp model to generate high-quality grasps for novel objects with different shapes and textures. Our work demonstrates the first deep, active learning approach to robotic grasping. Our real-robot grasping experiments show our active grasp planner achieves comparable success rates, while using fewer training data, when compared to a passive supervised planner [8, 13, 20] trained on data generated with a geometric grasp planner. Furthermore our active learner generates grasps with greater diversity than the passively trained model.

In the next section we review the literature of grasp planning and active learning in robotic grasping. In Section III we introduce our grasp planning as inference framework

<sup>1</sup>Qingkai Lu, Mark Van der Merwe, and Tucker Hermans are with the School of Computing and the Robotics Center, University of Utah, Salt Lake City, UT USA. <sup>2</sup>NVIDIA; Seattle, WA, USA. qklu@cs.utah.edu; mark.vandermerwe@utah.edu; thermans@cs.utah.edu. Q. Lu was supported by NSF Award #1846341 and DARPA Award under grant N66001-19-2-4035. M. Van der Merwe was supported in part by the Undergraduate Research Opportunities Program at the University of Utah.

and define our grasp model. We then describe our novel, bandit-based active grasp learning algorithm in Section IV. In Section V, we give a thorough account of our experimental evaluation. In Section VI, we conclude with a brief discussion and suggest directions for future work.

## II. RELATED WORK

Broadly speaking, supervised deep grasp learning methods either predict grasp success from an image patch associated with a grasp configuration [4, 5, 7, 9, 11, 14, 21, 22] or directly predict a grasp configuration from an image or image patch using regression [12, 23, 24]. Most deep grasp learning approaches such as [1, 4–7, 9, 25] focus on parallel jaw grippers, while relatively little work has focused on the more difficult multi-fingered grasping problem [8, 11–13, 26]. Deep learning requires large amount of training data, however, it is challenging to collect a large amount of grasp training data, especially for multi-fingered grasping. Representative approaches to grasp learning for parallel jaw grippers collect tens of thousands [5], hundred thousands [14], millions [7], or even one billion [15] grasp attempts.

Compared with grasp training data of two-fingered parallel jaw grippers, it is more difficult to collect grasping data for multi-fingered hand. Only 3,770 grasps were generated for the Barrett hand using GraspIt! [27] in [11], Lu et al. [8] collected only 1507 grasp attempts for the four-fingered Allegro hand, and Veres et al. [12] collected a data-set of around 47,000. In this paper, we propose a novel active grasp learning approach in order to address the data collection bottleneck for grasp learning on multi-fingered hands in order to avoid simply scaling up grasp attempts to significantly larger numbers.

Reinforcement learning stands as an alternative to supervised learning techniques for robotic grasping [25, 26]. However, the sample complexity of reinforcement learning generally requires even more data than supervised approaches as the robot must learn not only the necessary grasp configuration, but also how to reach and lift the target object. For example Kalashnikov et al. [25] had to collect over 580k grasps over the course of several weeks across 7 robots to learn to grasp with a two-fingered gripper. An alternative is to leverage human demonstrations to provide an initial supervised learning phase to the RL process, at the cost of giving up full autonomous exploration and labeling [28].

Active learning scenarios can be organized [16] as: 1) stream-based selective sampling where a stream of samples come in and the learner selects if the current observed sample should get a label; 2) pool-based sampling where the learner has a fixed set of unlabeled samples and can ask which sample can next get a label; and 3) query synthesis where the learner selects what input sample to generate and query for the label.

Kroemer et al. [17] combine active learning and reactive control for robot grasping. They propose a hierarchical controller to decide where and how to grasp the object. The upper

level controller selects where to grasp the object using a reinforcement learner, which models the reward distribution over the action space using Gaussian Process Regression (GPR). Active learning is applied to incorporate the supervised data into GPR. The lower level controller comprises an imitation learner and a vision-based reactive controller to determine appropriate grasping motions. Compared with our work, their grasp learning and planning are performed separately for each given object without sharing information across different objects. Moreover, they only model the 6 DOF hand pose without considering the hand joint angles. In [29], an active learning approach is presented for assessing robot grasp reliability. It classifies grasp success using the weighted  $k$ -nearest neighbors algorithm and defines a classification confidence metric for their pooling-based active learning. However, the grasp active learning in [29] is evaluated as an offline classification problem without performing grasp planning.

Montesano et al. [30] address the problem of actively learning good grasping points using a pooling formulation to reduce the number of examples needed for training the grasping model. This paper uses a non-parametric kernel approach to compute the success probability of grasping points. Its proposed active learning metric favors exploration of areas with high probability of success or where uncertainty is high. This pooling-based active learning is tested on a real humanoid robot for both seen and unseen objects. In [31], active perception is used to collect the training data set efficiently for each specific object, although no active learning is performed over grasp configuration or control. Tian et al. [32] compute an approximation of the grasp space for an object of interest using SVM-based active learning and apply bijective contact mapping to transfer grasp contact points from an example object to a novel object. However, their active learning focuses on the separation of in-collision grasp configurations from the collision free configurations, instead of the grasp quality.

Existing bandit-based grasping work [17–19] treats each possible grasp as one arm and perform grasp planning separately for different objects. Instead, we use the bandit-framework to select qualitatively different grasps to actively improve our grasp deep learning models in a unified manner across all objects in our training set. Moreover, none of these multi-armed bandit grasp planning work focus on rigid multi-fingered hand grasping like we do. In [17], active grasp learning for the Barrett hand is formulated as a continuum-armed bandits problem by treating each grasp pose as one arm. A continuum Gaussian Bandits algorithm is proposed that identifies the local maxima using a gradient-based method inspired by mean shift.

In [18, 33], multi-armed bandit is used to generate force closure grasps for parallel jaw grippers on each given 2D planar object. Mahler et al. [19] extends the multi-armed bandit grasp planning of [18] from 2D to 3D objects for parallel jaw grippers. The similarity between a pair of grasps

and objects is measured and used as the prior information for the multi-armed bandit grasp planning. The object feature is extracted using a Multi-View CNN. Thompson sampling is used to solve the grasping multi-armed bandit problem in [18, 19, 33]. Oberlin et al. [34] formalize grasping for parallel jaw grippers as a multi-armed bandit problem. It defines a new algorithm called Prior Confidence Bound for best arm identification in budgeted bandits, which enables the robot to quickly find an arm corresponding to a good grasp without pulling all the arms. A linear grasp model is used to score grasp candidates for their grasping multi-armed bandit. They take an instance-based approach that needs to collect training data for all testing objects to be grasped, which does not generalize to novel objects.

In [35], a contextual multi-armed bandit formulation selects one of three pre-defined environment-constrained grasping strategies for a given object, instead of learning the grasp configuration across different objects as we do. Different multi-armed bandit algorithms such as Thompson sampling and GP-UCB are compared.

Compared with existing grasp active learning work, our active learning approach has three novelties: the first grasp active learning work leveraging deep networks; doing a continuous synthesis approach instead of pooling for grasp active learning; modeling grasp active learning as a multi-armed bandit problem to cover the grasp space across different objects efficiently.

### III. GRASP PLANNING AS INFERENCE

Our grasping method has two stages. In the learning stage described in Section IV, we use active learning to train the grasp model comprised of the voxel-based classifier and the mixture density network (MDN) prior. In the planning stage, we perform maximum a posteriori (MAP) inference using this model to synthesize a grasp preshape configuration for execution on the robot.

In this work, we define the grasp configuration as the palm pose and the hand's preshape joint angles that define the shape of the hand prior to closing [13]. In order to make the grasp inference agnostic to object poses, we encode the palm pose in the object reference frame for learning. At inference time we optimize over the robot arm joint configuration solving the arm inverse kinematics jointly with the grasp configuration. After finding the grasp preshape configuration, the robot moves to this preshape and runs a controller to close the hand forming the grasp on the object. The specific joints that define the preshape can be found in [13].

We focus on scenarios where a single, isolated object of interest is present in the scene. Importantly, we assume no explicit knowledge of the object beyond a single camera sensor reading of it in its current pose. The problem we address is, given such a grasp scenario, plan a grasp preshape configuration that allows the robot to successfully grasp and lift the object without dropping it.

Given the learned model parameters,  $\mathbf{W}$  and  $\Phi$ , along with the visual representation,  $\mathbf{z}$ , associated with an observed

object of interest, our goal is to infer the grasp configuration  $\mathbf{q}$  in the robot arm joint configuration space that maximizes the posterior probability of grasp success  $Y = 1$ . Here  $Y$  defines a random Boolean variable with 0 meaning failure and 1 meaning success. We can thus formalize grasp planning as a *maximum a posteriori* (MAP) inference problem:

$$\begin{aligned} \underset{\mathbf{q}}{\operatorname{argmax}} \quad & p(\mathbf{q}|Y = 1, \mathbf{z}, \mathbf{W}, \Phi) \propto p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{W})p(\mathbf{q}|\mathbf{z}, \Phi) \\ \text{subject to} \quad & \mathbf{q}_{\min} \leq \mathbf{q} \leq \mathbf{q}_{\max} \end{aligned} \quad (1)$$

We constrain the grasp configuration parameters to obey the joint limits of the robot hand in Eq. (1). Other constraints such as collision avoidance could optionally be added [36].

We define the grasp success likelihood  $p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{W})$  to be a voxel-based 3D convolutional network following [13].  $\mathbf{W}$  represents the neural network parameters. The voxel-based network predicts the probability of grasp success,  $Y$ , as a function of the visual representation of the object of interest,  $\mathbf{z}$ , and grasp configuration,  $\mathbf{q}$ . Figure 1a shows the architecture of our grasp success prediction network.

We additionally learn a model of the conditional probability distribution  $p(\mathbf{q}|\mathbf{z}, \Phi)$  encoding the distribution of successful grasps given the object of interest. We construct this as a mixture-density network (MDN) following [13]. Given its input an MDN predicts the parameters (means, covariance and mixing weights) of a Gaussian mixture model as output.  $\Phi$  define the learned weights of the MDN. Figure 1b shows the architecture of our grasp prior. We train our voxel-based MDN using the negative log likelihood loss.

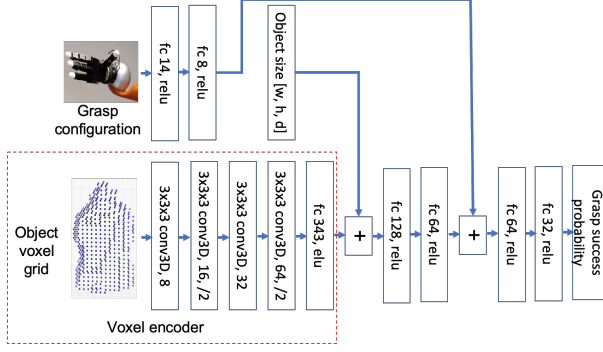
We use the object voxel-grid and the object size vector as the visual object representation [13]. In order to generate the voxel-grid we first segment the object from the 3D point cloud by fitting a plane to the table using RANSAC and extracting the points above the table. We then estimate the first and second principle axes of the segmented object to create a right-handed object reference frame aligned relative to the world frame. We compute the object size along the three coordinates of the object reference frame to construct the object size vector. We then generate a  $32 \times 32 \times 32$  voxel grid oriented about this reference frame. More details of the voxel-grid generation are described in [13].

We solve the grasp inference in the log-probability space and regularize the log-prior with a multiplicative gain of 0.5 to prevent the prior dominating the inference. We use the scikit-learn implementation of the popular L-BFGS optimization algorithm with bound constraints to efficiently solve the inference problem.<sup>1</sup> We initialize the optimization by sampling from the MDN prior.


### IV. ACTIVE GRASP LEARNING

As explained in Section II three primary active learning scenarios have been considered in the literature: stream-based selective sampling, pool-based sampling, and query

<sup>1</sup><http://scikit-learn.org/stable/index.html>



(a) The architecture of our grasp success classifier.



(b) The architecture of our grasp conditional prior.

**Fig. 1:** The voxel-config-net and MDN architectures [13]. Bottom left visualizes the voxel-grid for the “mustard bottle” object. All convolutional layers use  $3 \times 3 \times 3$  3D convolutional filters with exponential linear unit (ELU) activations. We annotate the number of filters and the stride (/2 means a stride of 2) for convolutional layers. We annotate the number of neurons and activation function for fully connected layers.

synthesis [16]. In this paper, we examine the query synthesis approach in a continuous fashion for active grasp learning.

Our active learning approach combines three separate exploration strategies: the likelihood uncertainty of the grasp success classifier, grasp success probability maximization, and the grasp configuration exploration. These three arms are designed to improve both the grasp success classifier and enable the grasp model to cover the space of grasp configuration as much as possible.

We formulate active learning as a multi-arm bandit problem to balance exploration and exploitation across our three different strategies. The multi-armed bandit is a classic reinforcement learning problem where we are given  $n$  slot machines (arms) [37]. Each slot machine has its own probability distribution of reward. At each time step, the agent chooses one slot machine to play and receives a reward. The agent’s objective is to decide which arm to play at each time step such that it can maximize the cumulative reward.

#### A. Active Grasp Learning Strategies

The **grasp likelihood uncertainty maximization arm** defines a standard active learning objective for binary classification problems in order to improve classification accuracy. We define the likelihood uncertainty maximization arm in Eq. (2).

$$\begin{aligned} & \underset{\mathbf{q}}{\operatorname{argmax}} \quad f(\mathbf{q}) + g(\mathbf{q}) \\ & \text{subject to} \quad \mathbf{q}_{\min} \leq \mathbf{q} \leq \mathbf{q}_{\max} \end{aligned} \quad (2)$$

**Algorithm 1:** Pseudocode of our multi-armed bandit active learning.

**Data:** The geometrical grasp data *GeoData* collected using a heuristic geometrical planner.

**Result:** The grasp model  $G$ , including the voxel-based classifier and the MDN conditional prior.

```

1 Initialize  $G$  through supervised learning using  $GeoData$ ;
2  $N = 128, M = 16, K = 4, E = 5$ ;
3  $i = 1$ ;
4 Initialize the active learning grasping data  $ActiveData$  to be an empty dataset;
  //  $N$ : Number of active learning rounds
5 while  $i \leq N$  do
6    $j = 0$ ;
7   Initialize  $CurRoundData$  to be an empty dataset;
  //  $M$ : Number of batches for an active learning round
8   while  $j < M$  do
9     Synthesize an active learning grasp query  $Q$  using the multi-arm bandit algorithm UCB;
10    Add  $Q$  into  $CurRoundData$ ;
11     $j = j + 1$ ;
12  end
13  Add  $CurRoundData$  into  $ActiveData$ ;
  //  $K$ : For every  $K$  rounds, online update using both  $GeoData$  and  $ActiveData$ 
14  if  $i \% K == 0$  then
15    Online update the grasp model  $G$  using both  $GeoData$  and  $ActiveData$  for  $E$  epochs;
16  else
17    Online update the grasp model  $G$  using  $CurRoundData$  for  $E$  epochs;
18  end
19   $i = i + 1$ ;
20 end

```

$$f(\mathbf{q}) = \frac{1}{2(1 + \exp(-\log p(\mathbf{q}|\mathbf{z}, \phi)))} \quad (3)$$

$$g(\mathbf{q}) = \begin{cases} p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{w}) & p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{w}) \leq 0.5 \\ 1 - p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{w}) & p(Y = 1|\mathbf{q}, \mathbf{z}, \mathbf{w}) > 0.5 \end{cases} \quad (4)$$

Given the learned model parameters  $\mathbf{W}$ ,  $\Phi$ , and the visual representation  $\mathbf{z}$  associated with an observed object of interest, our goal is to infer the grasp configuration parameters  $\mathbf{q}$  that maximize the uncertainty of the grasp success likelihood  $Y$ .  $f(\mathbf{q})$  in Eq. (3) represents the grasp success classification uncertainty.  $g(\mathbf{q})$ , defined in Eq. (4), regularizes the optimization to not stray into areas far from grasp configurations observed in the training data. We use L-BFGS with bound constraints to solve the uncertainty

maximization initialized with a configuration sampled from the prior. We treat its optimized objective function value as its bandit reward.

**The grasp success probability maximization arm** encourages the active learner to synthesize more successful grasps to overcome the common issue of class imbalance in multi-fingered grasp learning. For example in [8] only 11% out of 1507 training grasps generated using a geometric grasp planner were successful. We perform the same grasp inference as in Eq. (1) for this arm. We normalize the logarithm of the optimized grasp success posterior with a Sigmoid function as the reward for the bandit.

**The grasp configuration exploration arm** enables the grasp model to cover a larger portion of the grasp configuration space by exploring the areas farther from previous grasp attempts. To achieve this, we sample 50 grasp configuration candidates from the prior and select the one with lowest prior probability density as the grasp to explore. We compute the Sigmoid of the negative logarithm of the prior density as the reward of the selected grasp.

### B. Implementation Details

During training we randomly select objects to present to the active learner in order to cover the object space well. We solve the bandit problem using the classical Upper Confidence Bound (UCB) algorithm [37]. The UCB algorithm exploits actions with high average rewards obtained and explores more uncertain actions. Algorithm 1 shows the pseudocode of our multi-armed bandit active learning. Online training is a known difficult problem for active deep learning [38]. For each round of active learning, we first apply the multi-arm bandit algorithm to generate an active learning mini-batch that contains 16 grasps using the model of previous round (i.e. data acquisition). Then we online update the grasp model for each round, including the voxel-based classifier and the MDN prior. Depending on the active learning round number, we use either the mini-batch data of the current round or all the active and geometrical data for online training.

## V. EXPERIMENTS

In this section, we describe the experimental evaluation and analysis of our grasp active learner. We compare our active learning approach to passive supervised learning approach used to initialize our active learning grasp model. We then compare to the passive supervised learning model trained with more samples.

### A. Robotic System for Data Collection and Experiments

We conduct all training and experiments using the four-fingered, 16 DOF Allegro hand mounted on a Kuka LBR4 7 DOF arm. We evaluate our grasp planners on the physical robot. There are 15 parameters for the Allegro hand preshape, 7 for the LBR4 arm joint angles representing the palm pose and 8 relating to the first 2 joint angles of each finger proximal to the palm. We use a Kinect2 camera to generate

	Max success	Max uncertainty	Exploration
Average reward	0.965	0.922	0.933
Average time	8.11	4.15	2.53
Number of pulls	933	522	597

**TABLE I:** Average rewards and number of pulls for the 3 active learning arms.

the point cloud of the object on the table. One example RGB image of the robot and the object generated by Kinect2 can be seen from Figure 2.



**Fig. 2:** Example RGB image of the experimental robot setup from the RGB-D Kinect 2 camera.

### B. Active Learning Setup and Analysis


We train the initial grasp classifier and MDN prior, for active learning in a passive supervised way using the data from [13]. These data contain 10 sets of grasps on the Bigbird objects data set [39] containing a total of 8,988 grasps, 2308 of which were successful. We randomly select 5 sets (i.e. 4,578 grasp samples) out of the 10 training sets in [13] to train the initial grasp model for active learning. 1,168 out of these 4,578 grasps are successful grasps.

We run active learning for 128 rounds in simulation. Our active learner generates 16 grasp queries per round. We online update the grasp model for each round of active learning as described in Section IV. We randomly select a different Bigbird object every 5 active learning queries in order to cover the object space well.

The average rewards, average running time, and number of plays of our three active learning arms can be seen from Table I. We empirically add a constant of 0.35,  $-0.05$ , and 0.6 to the max success, max uncertainty, and exploration arm rewards, respectively, for the multi-armed bandit active learning. The constant terms cause the rewards of different arms to be in a similar range.

Figure 3 shows example grasps queried by the three different arms of our active learner. By measuring the differential entropy [40, 41] of the collected training data, we see that the active learner generates a larger diversity of grasp preshape configurations than the heuristic planner. The grasp palm pose for entropy computation is in Cartesian space, while hand joint angles are in radians. We generate 4 heuristic grasp subsets for entropy computation by random selection without replacement. Each subset has the same number of grasps as the active dataset. Table II summarizes these results. In Figure 4, we show preshape examples from the active






**Fig. 3:** Examples of grasp preshapes queries generated by our three active learning arms. The columns from left to right show grasp queries synthesized by the max success, max uncertainty, and the exploration arm respectively.

	Active learning	Heuristic mean	Heuristic std
Config entropy	-6.1	-33.6	0.05
Pose entropy	0.1	-1.5	0.05
Joint entropy	-6	-32.1	0.007

**TABLE II:** The differential entropy of multivariate Gaussian distributions fit to active learning and heuristic grasping data.

learner which the heuristic grasp planner would not able to generate.



**Fig. 4:** Active grasp preshape examples that the heuristic grasp planner during data collection is not able to generate.

### C. Real Robot Experiments

We compare our active learning model with the initial passive supervised model to examine if our active learner can improve on the initial model. We also compare our active learning model with a passively learned model trained on 8,988 grasps from the dataset of [13] using the same learning criteria as in that paper.

We evaluate grasp planning as inference using the active learning model, the passive supervised model trained with 8,988 geometrical grasps, and the passive supervised initial model trained with 4,578 geometrical grasps on the physical robot system. We perform experiments on 8 YCB [42] objects covering different textures, shapes, and sizes. We show the experimental setup and objects used in Figure 5. All experimental objects are unseen in training except for “Pringles”. We attempted grasps at 5 different poses per object, for a total of 40 grasp attempts per method. We use the same set of locations across different methods, but each object has its own set of random orientations. In total, we performed 120 grasp attempts for 3 different methods across 8 objects in this paper.

If the RRT-connect motion planner fails to generate a plan for a grasp due to collision avoidance, we generate a



**Fig. 5:** Experimental setup with objects used for experiments. From left to right objects are “pringles”, “pitcher”, “soft scrub”, “sugar box”, “mustard bottle”, “Lego”, “soccer ball”, and “mug”. Objects range in size from  $8 \times 9 \times 11$ cm (mug) to  $13 \times 17 \times 24$ cm (pitcher).


new grasp using the same grasp planner with a different initialization. If the grasp planner could not generate a grasp with a motion plan in 5 attempts, we treat the grasp attempt as a failure case. It turns out all three grasp planners can find a grasp with a feasible motion plan for every object pose we tested in 5 attempts for all objects but the mug where no method could find a successful side grasp.

As described in Section V-A, we label a grasp attempt that successfully lifts the object to a height of 0.15m without dropping it as successful. We also manually label each experimented grasp to be a side or overhead grasp.

The grasp success rates for all three methods are summarized in Figure 6. It takes around 3 – 10 seconds for each method to generate a grasp. The grasp planners using the active learning model, the passive supervised model with more training data, and the passive supervised model for initialization achieve grasping success rates of 50%, 52.5%, and 32.5% respectively for the 8 objects.

The grasp inference using the active learning model generated 15 side and 25 overhead grasps for the 8 test objects. The grasp inference using the passive supervised model with more training data plans 11 side and 29 overhead grasps, while the smaller data passive model generated 14 side and 26 overhead grasps. More overhead grasps than side grasps were generated for all three methods both in training and experiments. Overhead grasps are relatively further away from the table compared with side grasps, which makes it easier for the motion planner to avoid collision with the table.

We also report the grasp success rates of side and overhead grasps separately for each grasp planner in Figure 6. The grasp inference using the active learning model, the passive supervised model with more training data, and the passive supervised model for initialization achieve success rates of 93.3%, 100%, and 78.6% respectively for side grasps of the 8 objects. The grasp inference using the active learning model, the passive supervised model with more training data, and the passive supervised model for initialization achieve success



**Fig. 6:** Multi-fingered grasping success rates of grasp inference using 3 different models on the real robot. “Pringles” was seen in training, other 7 objects are previously unseen.



**Fig. 7:** Examples of successful grasps generated by inference with our actively learned grasp model. The top row shows side grasps. The bottom row shows overhead grasps.

rates of 24%, 34.5%, and 7.7% respectively for overhead grasps of the 8 objects.

All three grasp planners have lower success rate for overhead grasps than side grasps. Objects such as mustard, mug, and lego have relatively smaller contact areas available for overhead grasps and our grasp controller would push them away when closing the hand as it had no feedback from vision or haptic sensors to know the object was moving [13].

The grasp planner using the actively learned model outperforms the passive supervised model for initialization, which shows that active learning improves the supervised passive model for grasp inference. The grasp planner using the active learning model achieves comparable performance with the passive supervised learning with fewer grasp samples, which demonstrates the benefit of grasp active learning. This implies that our active grasp learning covers the grasp configuration space better across different objects with fewer samples, compared with passive supervised grasp learning. As noted above no planner could find a side grasp for the mug; additionally, all overhead grasp attempts for the mug failed during execution.

In Figure 7, example grasps are shown for different objects generated by our inference approach with the actively learned grasp model, including the voxel-based classifier and MDN prior. We show side grasps that provide stability in the top row. We present overhead grasps in the bottom row,

	Active	Supervised more	Supervised less
Config entropy	−16.3	−18.6	−15.9
Pose entropy	−3.7	−4.2	−3.7
Joint entropy	−11.8	−13.5	−11

**TABLE III:** The differential entropy of a multivariate Gaussian distribution fit to our real-robot experiment grasps for the three different models.

which offer dexterity and access to objects in clutter. Lastly, Table III shows the active learner plans real-robot grasps with greater diversity than the passively learned large-data model.

## VI. DISCUSSION AND CONCLUSION

In this work, we propose an approach to active multi-fingered, grasp learning using three different exploration strategies. We formalize these strategies as arms in a bandit setting. Our real-robot grasping experiment shows our active grasp planner trained using fewer data achieves comparable success rates with a passive supervised planner trained on data generated by a geometrical grasp planner. This implies our active grasp learning covers the grasp configuration space better across different objects with fewer samples, compared with passive supervised grasp learning. We also compute the differential entropy to demonstrate our active learner generates grasps with larger diversity than passive supervised learning using more heuristic data, which attains comparable success rate.

In the future, we will run more rounds of active learning in order to cover more of the grasping and objects space, which

can help to generate grasps with more diversities for different tasks. We plan to extend our approach to also actively select which object from a set to attempt a grasp on. We believe learning or designing a more complex feedback controller for overhead grasps using tactile feedback would boost the overhead grasp performance [13].

## REFERENCES

- [1] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-Driven Grasp Synthesis-a Survey," *IEEE Trans. on Robotics*, vol. 30, no. 2, pp. 289–309, 2014.
- [2] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic Grasping of Novel Objects using Vision," *Intl. Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.
- [3] A. Saxena, L. L. S. Wong, and A. Y. Ng, "Learning Grasp Strategies with Partial Shape Information," in *AAAI National Conf. on Artificial Intelligence*, 2008, pp. 1491–1494.
- [4] I. Lenz, H. Lee, and A. Saxena, "Deep Learning for Detecting Robotic Grasps," *Intl. Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [5] L. Pinto and A. Gupta, "Supersizing Self-supervision: Learning to Grasp from 50K Tries and 700 Robot Hours," in *IEEE Intl. Conf. on Robotics and Automation*, 2016, pp. 3406–3413.
- [6] M. Kopicki, R. Detry, M. Adjigble, R. Stolkin, A. Leonardis, and J. L. Wyatt, "One-Shot Learning and Generation of Dexterous Grasps for Novel Objects," *Intl. Journal of Robotics Research*, vol. 35, no. 8, pp. 959–976, 2016.
- [7] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics," in *Robotics: Science and Systems*, 2017.
- [8] Q. Lu, K. Chenna, B. Sundaralingam, and T. Hermans, "Planning Multi-Fingered Grasps as Probabilistic Inference in a Learned Deep Network," in *Intl. Symp. on Robotics Research*, 2017.
- [9] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," *arXiv preprint arXiv:1905.10520*, 2019.
- [10] A. Sahbani, S. El-Khoury, and P. Bidaud, "An Overview of 3D Object Grasp Synthesis Algorithms," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 326–336, 2012.
- [11] J. Varley, J. Weisz, J. Weiss, and P. Allen, "Generating Multi-Fingered Robotic Grasps via Deep Learning," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2015, pp. 4415–4420.
- [12] M. Veres, M. Moussa, and G. W. Taylor, "Modeling Grasp Motor Imagery through Deep Conditional Generative Models," *IEEE Robotics & Automation Letters*, vol. 2, no. 2, pp. 757–764, 2017.
- [13] Q. Lu, M. Van der Merwe, B. Sundaralingam, and T. Hermans, "Multi-fingered grasp planning via inference in deep neural networks," *IEEE Robotics & Automation Magazine*, 2020.
- [14] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection," *Intl. Journal of Robotics Research*, p. 0278364917710318, 2016.
- [15] C. Eppner, A. Mousavian, and D. Fox, "A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set," *arXiv preprint arXiv:1912.05604*, 2019.
- [16] B. Settles, *Active Learning*. Morgan & Claypool, 2012.
- [17] O. Kroemer, R. Detry, J. Piater, and J. Peters, "Combining active learning and reactive control for robot grasping," *Robotics and Autonomous Systems*, vol. 58, no. 9, pp. 1105–1116, 2010.
- [18] M. Laskey, J. Mahler, Z. McCarthy, F. T. Pokorny, S. Patil, J. Van Den Berg, D. Kragic, P. Abbeel, and K. Goldberg, "Multi-armed bandit models for 2d grasp planning with uncertainty," in *IEEE Intl. Conf. on Automation Science and Engineering*, 2015, pp. 572–579.
- [19] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. 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 *IEEE Intl. Conf. on Robotics and Automation*, 2016, pp. 1957–1964.
- [20] Q. Lu and T. Hermans, "Modeling Grasp Type Improves Learning-Based Grasp Planning," *IEEE Robotics & Automation Letters*, 2019.
- [21] M. Gualtieri, A. ten Pas, K. Saenko, and R. Platt, "High Precision Grasp Pose Detection in Dense Clutter," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2016, pp. 598–605.
- [22] E. Johns, S. Leutenegger, and A. J. Davison, "Deep Learning a Grasp Function for Grasping under Gripper Pose Uncertainty," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2016, pp. 4461–4468.
- [23] J. Redmon and A. Angelova, "Real-Time Grasp Detection using Convolutional Neural Networks," in *IEEE Intl. Conf. on Robotics and Automation*, 2015, pp. 1316–1322.
- [24] S. Kumra and C. Kanan, "Robotic Grasp Detection using Deep Convolutional Neural Networks," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2017.
- [25] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke *et al.*, "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation," *arXiv preprint arXiv:1806.10293*, 2018.
- [26] B. Wu, I. Akinola, A. Gupta, F. Xu, J. Varley, D. Watkins-Valls, and P. K. Allen, "Generative attention learning: a general framework for high-performance multi-fingered grasping in clutter," *Autonomous Robots*, pp. 1–20, 2020.
- [27] M. Ciocarlie, C. Goldfeder, and P. Allen, "Dimensionality Reduction for Hand-Independent Dexterous Robotic Grasping," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2007, pp. 3270–3275.
- [28] T. Osa, J. Peters, and G. Neumann, "Experiments with hierarchical reinforcement learning of multiple grasping policies," in *Int. Symp. on Exp. Robot.* Springer, 2016, pp. 160–172.
- [29] A. Morales, E. Chinellato, A. H. Fagg, and A. P. del Pobal, "An active learning approach for assessing robot grasp reliability," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2004, pp. 485–490.
- [30] L. Montesano and M. Lopes, "Active learning of visual descriptors for grasping using non-parametric smoothed beta distributions," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 452–462, 2012.
- [31] X. Fu, Y. Liu, and Z. Wang, "Active learning-based grasp for accurate industrial manipulation," *IEEE Transactions on Automation Science and Engineering*, 2019.
- [32] H. Tian, C. Wang, D. Manocha, and X. Zhang, "Transferring Grasp Configurations using Active Learning and Local Replanning," *arxiv*, 2018.
- [33] M. Laskey, Z. McCarthy, J. Mahler, F. T. Pokorny, S. Patil, J. Van Den Berg, D. Kragic, P. Abbeel, and K. Goldberg, "Budgeted multi-armed bandit models for sample-based grasp planning in the presence of uncertainty," in *IEEE Intl. Conf. on Robotics and Automation*, 2015.
- [34] J. Oberlin and S. Tellex, "Autonomously acquiring instance-based object models from experience," in *Intl. Symp. on Robotics Research*, 2015, pp. 73–90.
- [35] C. Eppner and O. Brock, "Visual detection of opportunities to exploit contact in grasping using contextual multi-armed bandits," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2017, pp. 273–278.
- [36] M. V. der Merwe, Q. Lu, B. Sundaralingam, M. Matak, and T. Hermans, "Learning Continuous 3D Reconstructions for Geometrically Aware Grasping," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020. [Online]. Available: <https://sites.google.com/view/reconstruction-grasp/home>
- [37] R. S. Sutton, A. G. Barto, and A. B. Book, *Reinforcement Learning : An Introduction*. MIT Press, 2018.
- [38] Y. Gal, R. Islam, and Z. Ghahramani, "Deep bayesian active learning with image data," in *Intl. Conf. on Machine Learning*, 2017.
- [39] A. Singh, J. Sha, K. S. Narayan, T. Achim, and P. Abbeel, "Bigbird: A Large-Scale 3D Database of Object Instances," in *IEEE Intl. Conf. on Robotics and Automation*, 2014, pp. 509–516.
- [40] L. Jost, "Entropy and diversity," *Oikos*, vol. 113, no. 2, pp. 363–375, 2006.
- [41] I. F. Spellerberg and P. J. Fedor, "A tribute to claude shannon (1916–2001) and a plea for more rigorous use of species richness, species diversity and the shannon–wienerindex," *Global ecology and biogeography*, vol. 12, no. 3, pp. 177–179, 2003.
- [42] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The YCB Object and Model Set: Towards Common Benchmarks for Manipulation Research," in *Intl. Conf. on Advanced Robotics*, 2015, pp. 510–517.