LEGS: Learning Efficient Grasp Sets for Exploratory Grasping

Letian Fu[1], Michael Danielczuk[1], Ashwin Balakrishna[1], Daniel S. Brown[1],
Jeffrey Ichnowski[1], Eugen Solowjow[2], Ken Goldberg[1]

Abstract— While deep learning has enabled significant progress in designing general purpose robot grasping systems, there remain objects which still pose challenges for these systems. Recent work on Exploratory Grasping has formalized the problem of systematically exploring grasps on these adversarial objects and explored a multi-armed bandit model for identifying high-quality grasps on each object stable pose. However, these systems are still limited to exploring a small number or grasps on each object. We present Learned Efficient Grasp Sets (LEGS), an algorithm that efficiently explores thousands of possible grasps by maintaining small active sets of promising grasps and determining when it can stop exploring the object with high confidence. Experiments suggest that LEGS can identify a high-quality grasp more efficiently than prior algorithms which do not use active sets. In simulation experiments, we measure the gap between the success probability of the best grasp identified by LEGS, baselines, and the most-robust grasp (verified ground truth). After 3000 exploration steps, LEGS outperforms baseline algorithms on 10/14 and 25/39 objects on the Dex-Net Adversarial and EGAD! datasets respectively. We then evaluate LEGS in physical experiments; trials on 3 challenging objects suggest that LEGS converges to high-performing grasps significantly faster than baselines. See `https://sites.google.com/view/legs-exp-grasping` for supplemental material and videos.

I. INTRODUCTION

Recent advances in deep learning have enabled the development of universal grasping systems that can robustly grasp a wide variety of objects [23–25, 28, 29, 34]. However, these systems can still struggle to grasp objects with adversarial [27, 35] geometries or which are significantly out of distribution from the objects seen during training. This problem is common in many industrial settings, in which newly manufactured machine parts for custom applications may look very different from the objects in the datasets typically used for training universal grasping systems.

Recently, bandit-style algorithms have been used to augment general-purpose grasping policies by rapidly adapting them to specific objects [11, 19, 21, 22]. Recently, Danielczuk *et al.* [8] introduced Exploratory Grasping, where a robot learns to grasp novel objects through online exploration of grasps and stable poses. Their algorithm, Bandits for Online Rapid Grasp Exploration Strategy (BORGES), learns robust pose-specific grasping policies. However, BORGES limits exploration to a fixed set of 100 grasps per stable pose, possibly to overlooking other high-quality grasps.

In this work, we extend Danielczuk *et al.* [8] to explore thousands of grasps per stable pose. Considering grasp sets of this scale increases the likelihood of converging to a

[1]The AUTOLab at UC Berkeley, [2]Siemens Research Lab

Fig. 1: **Top: LEGS in Physical Experiments:** LEGS repeatedly attempts grasps on an object, and if the grasps are successful, it re-drops the object into a new stable pose. **Bottom: LEGS Active Set Evolution:** LEGS works by adaptively curating a small active set of promising grasps out of a large reservoir of grasp candidates (left). As exploration progresses, LEGS refines its active set (shown in bolded red/green) to contain higher quality grasps (right).

robust grasp, but also makes efficient exploration challenging. To address this challenge, we propose *Learned Efficient Grasp Sets* (LEGS), which adaptively curates an active set of promising grasps rather than restricting exploration to a small fixed subset. The key insight is to use a combination of priors from a universal grasping system and online trials to maintain confidence bounds on grasp-success probabilities. LEGS uses these bounds to (1) update the grasps in its active set and (2) decide when to stop exploring.

This paper makes the following contributions: (1) a novel adaptive multi-armed bandits algorithm that curates a small set of high-performing grasps by actively removing and resampling grasps based on performance bounds and a novel termination condition that enables a robot to predict (with high confidence) when it reaches a desired level of performance; (2) a *self-supervised* physical grasping system where a robot explores candidate grasps with minimal human intervention (roughly 1 in every 100 grasp attempts); (3) simulation and physical experiments suggesting that LEGS can identify higher quality grasps within a fixed time horizon than prior algorithms which do not learn an active set.

II. RELATED WORK

A. Universal Grasping Algorithms

Recent robotic grasping algorithms generalize to a wide range of objects [18]. Open-loop algorithms synthesize grasps and predict their quality based on the geometry of the object, and then plan and execute a motion to attempt a high-quality grasp without feedback [20, 23–25, 29]. Closed-loop grasp planners that use vision-based gripper servoing [28, 34] and RL [15, 16] have also been popular in prior work. LEGS

is designed to leverage priors from these universal grasping algorithms to efficiently learn a robust grasp policy for a specific, difficult-to-grasp object [27, 35]. We use priors from Dex-Net 4.0 [25], a general grasp planner that learns a grasp-quality estimator from a large dataset of 3D object models in simulation and then uses this estimator to sample and evaluate the quality of grasps in physical trials.

B. Multi-Armed Bandits

Prior work on multi-armed bandits [31] has studied settings where the number of actions is large compared to the number of timesteps allocated for exploration [2, 6, 13, 14, 33, 36, 37]. One popular algorithmic framework for this setting is called *best arm identification*, where the goal is to adaptively reject a set of arms from consideration when there is high confidence that they are suboptimal [1, 5, 17]. LEGS builds on these ideas, by adaptively filtering actions from an active set by maintaining confidence bounds on the reward corresponding to each action. This mechanism makes it possible to efficiently perform best arm identification across multiple bandits problems, where each bandit problem represents a distinct stable pose of an object. LEGS can quickly converge to high-quality grasps on problems with thousands of grasps per stable pose.

C. Exploratory Grasping

Universal grasping algorithms often struggle with certain objects [27, 35]. Danielczuk *et al.* [8] show that grasping algorithms such as Dex-Net [25] are difficult to fine-tune online on such objects, and propose *Exploratory Grasping*, a problem formulation where the objective is to perform rapid online adaptation to grasp specific, unknown objects. To achieve this, prior works sample a fixed set of grasps on specific object stable poses and apply multi-armed bandit algorithms to rapidly identify high-performing candidates [11, 19, 21, 22]. Danielczuk *et al.* [8] extend these ideas with BORGES, which explores grasps across all object stable poses by using Thompson sampling and a learned Dex-Net prior [21]. However, BORGES can often overlook high-quality grasps since it restricts exploration to a small initial set of grasps. To address this issue, LEGS begins with a large set of grasp candidates and adaptively curates sets of promising grasps by adding and removing grasp candidates during exploration. By doing this, LEGS is able to converge to better long-term performance than BORGES (which uses a small fixed set of grasps), while also learning to robustly grasp an object faster than baselines that seek to directly explore large sets of grasp candidates.

III. PROBLEM STATEMENT

Overview: Given a difficult-to-grasp polyhedral object of unknown geometry that rests on a planar surface and is viewed by an overhead depth camera, we seek to learn to successfully grasp the object in all of its stable poses.

Problem Setup: Given a polyhedral object o, let N be its number of stable poses. Each stable pose $s \in \{1, 2, \ldots N\}$ is associated with a landing probability λ_s, which indicates the

probability of the object landing in pose s when released from sufficient height in a randomized orientation [12, 26]. Following Danielczuk *et al.* [8], we model our problem as a finite-horizon Markov Decision Process $\mathcal{M} = (\mathcal{S}, \mathcal{A}, T, R, H)$. We let $\mathcal{S}$ be the set of equivalence classes of distinguishable stable poses of the object and $\mathcal{A}$ be the set of all possible grasps on the object. Thus, $\mathcal{A} = \bigcup_{s \in \mathcal{S}}^{N} \mathcal{A}_s$, where $\mathcal{A}_s$ are the grasps available at a stable pose s. Given a grasp action a in stable pose s, the transition function $T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ determines the probability distribution over next stable poses. The reward function $R : \mathcal{S} \times \mathcal{A} \to \{0, 1\}$ is binary: a grasp is successful and $R(s, a) = 1$ if the grasped object does not fall from the gripper after it is lifted, and $R(s, a) = 0$ otherwise. Let $p_{s_a} = \mathbb{E}[R(s, a)]$ be the expected success probability of grasp a on stable pose s. We define a grasping policy as: $\pi : \mathcal{S} \times \mathcal{A} \to [0, 1]$, where $\pi(a|s)$ denotes the probability of selecting grasp a in pose s. We denote the finite horizon of the MDP as H. The robot initially does not know any of the stable poses or the number of stable poses N. If a grasp is successful, the robot randomizes the orientation of the object in the gripper, drops the object so that the next stable pose s' is determined by the landing probabilities $\{\lambda_s\}_{s=1}^{N}$, and records the observed stable pose s'.

We represent the actions, $\mathcal{A}_s$, at each stable pose s as candidate grasps sampled on the object. We use the same method as Mahler *et al.* [25] to sample antipodal grasps on each stable pose. We do not make any assumptions on the grasping modality, so in practice these grasps can be sampled from various different grasp planners, including parallel-jaw or suction grasp planners. We denote the number of possible grasps for pose s as $K_s = |\mathcal{A}_s|$ and the total number of grasps over all states as $K = \sum_{s \in \mathcal{S}} K_s$.

An important difference between our problem setting and prior work [8] is that we consider settings in which K is large (> 1000) and thus is of the same order of magnitude as the exploration horizon, H. This significantly exacerbates exploration challenges, since there is not enough time to fully explore each grasp, motivating the key innovations in LEGS.

Assumptions: In this work, we assume access to the following: (1) a grasp sampler which accepts as input a depth map and outputs a set of candidate grasp configurations on the surface of the depth map with associated robustness values; (2) a robot/gripper that can either execute these grasps or detect that they are in collision; (3) sufficient information in the camera image to detect whether the object stable pose changes; (4) an evaluation function to detect whether a grasp is successful. We note that these assumptions are satisfied by the system we build to instantiate LEGS in practice. In addition, we make the following assumptions about object's interaction with the environment: (5) if a grasp is unsuccessful, the object either remains in the same stable pose or topples into another stable pose; and (6) there exists a grasp with non-zero success probability on each stable pose. These last two assumptions are consistent with [8].

Metrics: We define the *optimality gap*, Δ_π as

$$\Delta_\pi = \mathbb{E}_{s \in \mathcal{S}} \left[p_s^* - p_{s_{\pi(s)}} \right] = \sum_{s \in \mathcal{S}} \lambda_s \cdot \left(p_s^* - p_{s_{\pi(s)}} \right), \quad \text{(III.1)}$$

where $p_s^* = \max_{a \in \mathcal{A}_s} \mathbb{E}[R(s,a)]$ and $p_{s_{\pi(s)}} = \mathbb{E}[R(s, \pi(s))]$. In simulation, we can evaluate the ground-truth grasp-success probability for a given grasp with robust quasi-static grasp wrench space analysis [38]. We thus approximate p_s^* by sampling a large number of grasps on each stable pose. Intuitively, the optimality gap Δ_π measures the expected difference, across all stable poses, between the optimal policy, which selects the best available grasp, and the policy π. In physical experiments, the optimality gap cannot be computed so we report the grasp-success rate of the learned policy π.

The objective is to find a policy that minimizes the optimality gap for a given object within H grasp attempts. Denoting a policy learned after H grasp attempts by π_H, the objective is to identify π_H^* such that:

$$\pi_H^* = \arg\min_{\pi_H} \Delta_{\pi_H}. \quad \text{(III.2)}$$

IV. Learned Efficient Grasp Sets

We propose Learned Efficient Grasp Sets (LEGS), a multi-arm bandits algorithm that uses confidence bounds on grasp-success probability to maintain a small active set of candidate grasps. LEGS starts with an estimate of the prior success probabilities for all grasps in a large reservoir of possible grasps, and updates their grasp-success probabilities based on online grasp trials using Thompson sampling as in Danielczuk *et al.* [8]. However, unlike BORGES, LEGS uses the priors and online grasp trials to construct confidence bounds on the grasp-success probabilities for each grasp (Section IV-A).

LEGS is summarized in Algorithm 1. Once LEGS visits a stable pose s, it checks whether it has visited s (line 4). In Sec. VI, we describe how to recognize stable poses in the physical setup. If the stable pose s has never been visited (line 5), LEGS adds the stable pose to the set of visited stable poses $\hat{S}$ (line 6) and initializes an active set of candidate grasps, $\tilde{A}_s$, along with the parameters of a Beta distribution associated with each grasp in the active set (lines 7-8). We rank the grasps in the reservoir by their estimated grasp success probabilities under the Grasp Quality Convolutional Neural Network (GQ-CNN) from Dex-Net 4.0 [25] and select the $k = 100$ grasps with the highest values. In each iteration, LEGS executes the grasp with the highest sampled value from the posterior (lines 9-11), observes the outcome (line 12), and updates the posterior distribution [30] (lines 13-16). In conjunction, LEGS also constructs confidence bounds on each of the success probabilities of each grasp (Section IV-A). Every n iterations, it uses these confidence bounds to identify and remove the grasps with low robustness (Section IV-B) (line 18), and replaces them with newly sampled grasps where grasps are ranked by their estimated grasp success probabilities under GQ-CNN (lines 19-20).

A. Constructing Confidence Bounds on Robustness

To determine which grasps to remove from the active set, LEGS constructs upper and lower confidence bounds on

---

Algorithm 1: Learned Efficient Grasp Sets (LEGS)

1 **Input:** object o, grasp sampler f_θ, resample interval n, number of active grasps k
2 Initialize the set of visited stable poses $\hat{S} = \varnothing$
3 **for** $t = 1, 2, \ldots$ **do**
4   Recognize the current stable pose s
5   **if** $s \notin \hat{S}$ **then**
6      $\hat{S} \leftarrow \hat{S} \cup \{s\}$
7   Use f_θ to sample k grasps as the active set $\tilde{A}_s$
8   Set α_i and β_i based on prior for all $i \in \tilde{A}_s$
9   **foreach** *grasp* $i \in \tilde{A}_s$ **do**
10   sample $\phi_i \sim Beta(\alpha_i + 1, \beta_i + 1)$
11   Execute grasp $i = \text{argmax}_{j \in \tilde{A}_s} \phi_j$
12   Observe reward $r = R(s, i)$
13   **if** $r = 1$ **then**
14      $\alpha_i \leftarrow \alpha_i + 1$
15   **else**
16      $\beta_i \leftarrow \beta_i + 1$
17   **if** $t \equiv 0 \pmod{n}$ **then**
18   Remove the grasps in $\mathcal{B} = \mathcal{B}_\ell \cup \mathcal{B}_\gamma$ from $\tilde{A}_s$ (see equations (IV.1) and (IV.2))
19   Sample $|\mathcal{B}|$ new grasps using f_θ
20   For each new grasp $j = 1, \ldots |\mathcal{B}|$, set α_j, β_j using prior from f_θ and add new grasp to $\tilde{A}_s$

---

grasp robustness. We model the success probability of grasp i via $X_i \sim Beta(\alpha_i, \beta_i)$, and empirically select a confidence threshold δ. Then the percent-point function $\text{PPF}(X_i, \delta)$, the inverse of the cumulative distribution function $F_{X_i}(x)$, returns the value x such that $F_{X_i}(x) = \delta$. The $(1 - \delta)$-lower and -upper confidence bounds for X_i are $X_{i,\ell} = \text{PPF}(X_i, \delta)$ and $X_{i,u} = \text{PPF}(X_i, 1 - \delta)$, respectively. As a grasp is sampled more often, the interval $[X_{i,\ell}, X_{i,u}]$ tightens to reflect increased certainty in the robustness of the grasp.

B. Posterior Dependent Grasp Removal

LEGS avoids over-exploring less robust grasps by identifying and removing grasps from the active set that are highly likely to be either (1) inferior to another grasp in the active set (*locally suboptimal*) or (2) below a desired global grasp success probability threshold (*globally suboptimal*). Let the highest lower confidence bound across all active grasps be: $X_\ell^* = \max_{i \in \tilde{A}_s} X_{i,\ell}$. We define the set of *locally suboptimal grasps* as the set of grasps for which their $(1-\delta)$-confidence upper bound is worse than the $(1-\delta)$-confidence lower bound for the best grasp in the active set:

$$\mathcal{B}_\ell = \{i : X_{i,u} < X_\ell^*\}. \quad \text{(IV.1)}$$

Thus, $\mathcal{B}_\ell$ represents the set of grasps that are likely to be inferior to the best known grasp in the active set. However, in the early stages of exploration, we may not yet have sampled a high-performing grasp and $\mathcal{B}_\ell$ may be empty. In these cases, we still desire to remove and resample grasps that, with high-confidence, are clearly low performing. Thus, given a minimum performance threshold $\gamma \in [0, 1]$, we define the set of *globally suboptimal* grasps in the active set (denoted

$\mathcal{B}_\gamma$): grasps which have been sampled, but are likely to have success probability less than γ. We define $\mathcal{B}_\gamma$ as

$$\mathcal{B}_\gamma = \{i : X_{i,u} < \gamma\}. \qquad (\text{IV.2})$$

We denote the set of attempted grasps in the active set as $\mathcal{P}$, and let the index of the currently known best grasp be $i*$. The full set of grasps removed by LEGS is constructed by taking the union of the above sets: $\mathcal{B} = (\mathcal{B}_\ell \cup \mathcal{B}_\gamma) \cap \mathcal{P} \setminus \{i*\}$. This allows LEGS to remove grasps which are unlikely to outperform the best known grasp in the current active set.

C. Early Stopping

Rather than setting the exploration horizon H to a fixed value, we can set a performance threshold and let LEGS stop exploring once it has high confidence that it has achieved the desired threshold. This early stopping condition allows LEGS to efficiently allocate exploration time by only continuing to explore objects that it cannot yet robustly grasp.

Given a user-specified, minimum performance threshold $\rho_{\min} \in [0,1]$, we want to detect when, with high likelihood, the true performance of LEGS is above this threshold. More formally, given a confidence parameter $\delta_{\text{stop}} \in [0,1]$, we want to calculate a $(1-\delta_{\text{stop}})$-confidence lower bound, denoted by p_ℓ, on the true expected performance of the grasping policy π, i.e., we want to find p_ℓ such that $Pr\left(p_\ell \leq \mathbb{E}_{s \in S}[p_{s_{\pi(s)}}]\right) \geq 1 - \delta_{\text{stop}}$. Then, the robot can stop exploring when $p_\ell \geq \rho_{\min}$.

We cannot directly compute $\mathbb{E}_{s \in \mathcal{S}}[p_{s_{\pi(s)}}]$ since we do not know the true stable pose distribution $\mathcal{S}$. Thus, we take a Bayesian approach where we approximate p_ℓ by sampling likely values of $\mathbb{E}_{s \in S}[p_{s_{\pi(s)}}]$ given the observed data and then by taking the δ_{stop}-percentile of these samples [3, 4]. First, for each observed stable pose, s, we estimate the expected performance of the best grasp as $\hat{p}_s^* = \max_{i \in \mathcal{A}_s} \frac{\alpha_i}{\alpha_i + \beta_i}$, where α_i and β_i are the parameters of the Beta posterior distribution over the success probability of grasp i. To reason about the performance of LEGS, we must account for uncertainty over the stable pose distribution, parametrized by the drop probabilities $\lambda_1, \ldots, \lambda_N$. However, N is unknown. Thus, we model our belief over drop probabilities using a Dirichlet posterior distribution over $\hat{N}+1$ drop probabilities, where $\hat{N}$ is the number of observed stable poses and the +1 allocates probability mass to unobserved stable poses.

Assuming a uniform Dirichlet prior, we take the empirical drop counts $c_1, \ldots, c_{\hat{N}}$ for $\hat{N}$ observed stable poses, and sample from the posterior distribution over stable pose drop probabilities, $Pr(\{\lambda_s\}_{s=1}^{\hat{N}+1} \mid c_1, \ldots, c_{\hat{N}}, 0)$. Due to conjugacy [10], the desired posterior distribution is also a Dirichlet distribution with parameters $(\alpha_1 = c_1 + 1, \ldots, \alpha_{\hat{N}} = c_{\hat{N}} + 1, \alpha_{\hat{N}+1} = 1)$. Given a sample, $\{\lambda_s'\}_{s=1}^{\hat{N}+1}$, from the above Dirichlet posterior, we transform it into a sample from the posterior over expected grasp robustness: $p_\pi' = \sum_{s=1}^{\hat{N}} \hat{p}_s^* \cdot \lambda_s'$. where we conservatively assume that the robot will fail to grasp the object in any unseen poses. We calculate a $(1-\delta_{\text{stop}})$-confidence lower bound on the overall grasp robustness by finding the δ_{stop} percentile, $\hat{p}_\ell = \text{PPF}(p_\pi', \delta_{\text{stop}})$, using M samples of p_π'.

V. SIMULATION EXPERIMENTS

A. Experimental Setup

We first evaluate LEGS in Exploratory Grasping with a variety of adversarial objects in simulation. Same as in Danielczuk et al. [8], we consider 14 Dex-Net 2.0 Adversarial objects [24] and all 39 EGAD! Adversarial evaluation objects [27]. We use Dex-Net 4.0 [25] to sample a large reservoir of $K = 2000$ grasps for each stable pose. We also use GQ-CNN to set the Beta prior for LEGS following the method from [8, 21]. Using the method outlined in Section IV, we update the active grasp set after every $n = 100$ timesteps and use $\delta = 0.05$ for constructing grasp confidence intervals with upper confidence threshold $\gamma = 0.2$. All experiments use a time horizon of $H = 3000$. We run 10 trials of each algorithm with 10 rollouts per trial, where each trial involves sampling a different reservoir of grasps, and each rollout for a trial involves running a grasp exploration algorithm.

B. Baselines

We compare LEGS against five baseline algorithms: Dex-Net, Tabular Q-Learning (TQL), BORGES ($K_s = 100$), BORGES ($K_s = 2000$), and LEGS (-AS). Dex-Net greedily chooses the best grasp evaluated by Dex-Net 4.0 [25] for each stable pose and does not do any online exploration. BORGES ($K_s = 100$) leverages a prior calculated by GQ-CNN to seed grasp success probability estimates, and then performs Thompson Sampling for each encountered stable pose to explore an initial active set of 100 grasps sampled on each of the poses. While BORGES ($K_s = 100$) is provided with the same initial active set as LEGS, unlike LEGS, BORGES ($K_s = 100$) does not update its set over time. However, different from [8], it is not guaranteed that there will exist successful grasps on all stable poses when $K_s = 100$. This implies that BORGES ($K_s = 100$) may not be able to transit between stable poses. The $K_s = 100$ Upper Bound refers to the optimality gap if on each stable pose, the best grasp in the active set is selected. BORGES ($K_s = 2000$) is identical to BORGES ($K_s = 100$), but instead directly explores the full reservoir of $K_s = 2000$ sampled grasps. TQL implements tabular Q-learning on the full reservoir of $K_s = 2000$ sampled grasps where each pose is a separate state s and each action a is a grasp on that pose and a Q-table $Q[s,a]$ is constructed to keep track of the corresponding 1-step Q-values. The values in the Q-table are initialized using the GQ-CNN prior and actions are chosen based on an ϵ-greedy policy [32] with $\epsilon = 0.1$. Finally, LEGS (-AS) is not provided with an initial active set, but instead operates on the full reservoir of $K_s = 2000$ grasps and uses the posterior dependent removal procedure in Section IV-B to remove grasps from the reservoir.

C. Experimental Results

We first study aggregated results of LEGS and baselines over objects in the Dex-Net Adversarial and EGAD! evaluation datasets in Table I. We find that LEGS performs better than or equal to the baseline algorithms on 10 out of 14 objects in the Dex-Net Adversarial dataset, and on 25 out of 39 objects in the EGAD! evaluation dataset. In

Dataset	Dex-Net	TQL	BORGES $(K_s = 100)$	$K_s = 100$ Upper Bound	BORGES $(K_s = 2000)$	LEGS (-AS)	LEGS
Dex-Net	0.56 ± 0.07	0.23 ± 0.08	0.13 ± 0.07	0.08 ± 0.04	$\mathbf{0.04 \pm 0.03}$	0.22 ± 0.06	$\mathbf{0.04 \pm 0.03}$
EGAD!	0.59 ± 0.03	0.32 ± 0.04	0.25 ± 0.04	0.13 ± 0.03	0.17 ± 0.03	0.28 ± 0.04	$\mathbf{0.14 \pm 0.03}$

TABLE I: **Grasping in Simulation Aggregated Results:** We show the optimality gap (mean $\pm$ standard error) achieved by LEGS and baselines after $H = 3000$ steps of exploration averaged over the objects in the Dex-Net Adversarial and EGAD! evaluation datasets. LEGS achieves a lower optimality gap than all baselines, indicating that LEGS is able to discover new high-performing grasps.
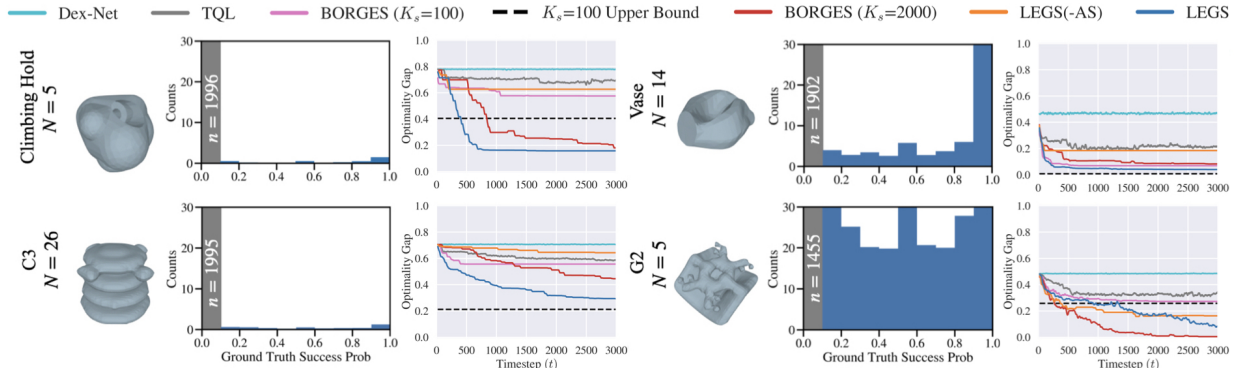

Fig. 2: **Simulated Grasping Experiments Case Study:** We report the performance of LEGS and baselines on four specific objects to investigate how object properties affect performance. For each object, we include a 3D rendering of the object and the number N of stable poses (left), a histogram of the ground truth grasp success probabilities over 2000 sampled grasps (middle), and learning curves (right).

comparison, the best performing baseline algorithm, BORGES $(K_s = 2000)$, only performs at least as well as rest of the algorithms on 5 out of 14 Dex-Net Adversarial objects and 14 out of 39 EGAD! evaluation dataset. On all of these objects we find that Dex-Net, which is not updated online, has high optimality gap, motivating online grasp exploration. The improvement for LEGS over LEGS (-AS) and BORGES $(K_s = 2000)$ indicates the increased efficiency of restricting exploration to a small active set, while the gap between LEGS and BORGES $(K_s = 100)$ indicates the importance of updating this active set over time to prune poor performing grasps while discovering new, high-quality grasps outside of the initial active set. BORGES $(K_s = 100)$ cannot outperform the success rate of the best grasp in its initial set $(K_s = 100$ upper bound). By contrast, LEGS, retains the efficiency of only exploring a small set of grasps while also being able to adapt this set over time to obtain successful grasps on difficult-to-grasp stable poses and reach a lower optimality gap. TQL learns much more slowly than BORGES because it fails to leverage the structure in the grasp exploration problem and does not learn separate policies for each stable pose.

In Figure 2, we study LEGS and baselines on specific objects. We show two objects (Climbing Hold and C3) where LEGS converges faster to high performing grasps than prior algorithms and two objects (F6 and Turbine Housing) where LEGS does not outperform all baselines. We find that when high performing grasps are abundant, LEGS may converge to suboptimal grasps. However, when there are only few successful grasps, LEGS can converge to good grasps much faster than baselines. If high quality grasps are already in the active set, LEGS can rapidly distinguish them from other grasps. If the active set does not contain successful grasps, LEGS can quickly replace bad grasps in the active set.

D. Early Stopping Results

Next, we study the accuracy and effectiveness of the early stopping criterion (Section IV-C). We test the proposed high-confidence performance bound across all objects in the Dex-Net Adversarial object set (individual results per object are reported in the supplement). We check whether LEGS has reached the stopping condition every 100 grasps for a horizon of $H = 3000$ total grasp attempts and use $\delta_{\text{stop}} = 0.05$, resulting in a 95%-confidence lower bound $\hat{p}_\ell$. We sample $M = 3000$ samples to estimate $\hat{p}_\ell$.

We first test how often the predicted bound is a true lower bound on performance. We find that, on average, across all Dex-Net Adversarial objects, our empirical lower bound is a 95.8%-accurate lower bound on the true performance over the true stable pose distribution. Thus, $\hat{p}_\ell$ forms an empirically valid $(1 - \delta_{\text{stop}})$-confidence lower bound. We next test the tightness of our lower bound. On average, the difference between the true performance of LEGS and our empirical lower bound is only 2.97%. These results suggest that our lower bound is highly accurate and tight enough to provide a practical signal for when the robot can safely stop exploring.

We next study, in simulation, the use our high-confidence bounds on performance for early stopping. As described in Section IV-C, given a user-specified, minimum performance threshold, ρ_{min}, the robot stops exploring when the lower confidence bound $\hat{p}_\ell$ is greater than ρ_{min}. When the robot chooses to stop exploring the object, we evaluate the ground truth performance of the learned policy and evaluate whether the true performance is also above the threshold ρ_{min}. We evaluate a wide range of thresholds and plot the results in Figure 3. Results suggest that we can achieve highly accurate early stopping, allowing the robot to accurately terminate exploration well before the full horizon of 3000 steps.

Fig. 3: **Early Stopping Threshold Sensitivity:** We evaluate early stopping over the Dex-Net Adversarial object set in simulation with a range of stopping thresholds, $\rho_{\min}$. We use a 95%-confidence lower bound on expected grasp robustness. **Left:** We plot the accuracy averaged over all objects and find that our empirical lower bound (Section V-D) is highly accurate across all stopping thresholds, $\rho_{\min}$. **Right:** We plot the number of steps before stopping, averaged across all objects. Intuitively, the required exploration time increases with higher performance thresholds. Importantly, the average number of steps before stopping is much lower than the 3000-step horizon.
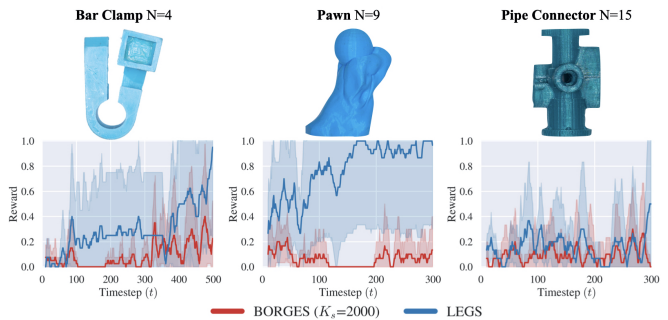

Fig. 4: **Physical Experiments Results:** We compare LEGS with BORGES ($K_s = 2000$) on three objects (Bar Clamp, Pawn, and Pipe Connector) from the Dex-Net Adversarial Dataset [24] in physical experiments. All physical experiments are completed within 3 hours. LEGS significantly outperforms BORGES ($K_s = 2000$) on Bar Clamp and Pawn, with minor improvement on Pipe Connector.

VI. PHYSICAL EXPERIMENTS

In this section, we discuss our experimental setup for physical experiments, the methods we used to enable intervention-free grasp exploration on a physical robot and results evaluating the performance of LEGS and BORGES ($K_s = 2000$) across 3 physical objects.

A. Experimental Setup

To deploy exploratory grasping algorithms on a physical robot, we modify the perception system introduced in Danielczuk *et al.* [8] to sample grasps and identify changes in the object stable pose. We capture a depth image of the object from an overhead camera, deproject it into a point cloud using the known camera intrinsics, demean the point cloud, and apply 3600 evenly spaced rotations to the point cloud around the camera's optical axis. We measure the chamfer distance between the rotated point clouds with previously cached point clouds and find the pair of point clouds that serves as the closest match. As in Danielczuk *et al.* [8], if at least 80 % of the points are less than 0.02 mm away from the closest points in the cached point cloud, we classify the two point clouds as belonging to the same stable pose. If none of the cached point clouds satisfies this condition, the point cloud is cached and treated as a new stable pose. If there exists a matching point cloud, we further align the translation and rotation of the point cloud via iterative closest point [7].

Upon discovery of a new stable pose, we use Dex-Net 4.0 [25] to sample, evaluate, and cache grasps in the grasp reservoir. Thus, LEGS can explore grasps on objects with unknown geometries and unknown numbers of stable poses.

B. Self-Supervised Exploratory Grasping with LEGS

Danielczuk *et al.* [8] find that re-dropping the object during experiments often cause it to fall out of the workspace, requiring extensive human effort to reset the object. To enable the robot to collect grasp data without human intervention, we introduce strategies to prevent the object from toppling out of the workspace while maintaining access to a wide variety of grasps. We drop the object within a bowl (Fig.1), where the object's rebound height is lower than the rim of the bowl.

The bowl allows the object to stay in the visible range of the overhead camera. However, the bowl's rim can be an obstacle to grasps. We introduce two autonomous *reset* behaviors to address this: (1) we center the object above the bowl before dropping the object, ad (2) when the object topples near the boundary, the robot pushes the object towards the center of the bowl to improve grasp access [9].

C. Experimental Results

Figure 4 shows learning curves from physical experiments comparing LEGS with BORGES ($K_s = 2000$) on three challenging objects from the Dex-Net Adversarial Dataset [24]. We run 3 trials with 1 rollout per trial for each object. We find that on 2 out of the 3 objects, LEGS is able to outperform BORGES ($K_s = 2000$) and identify high-performing grasps within a few hundred timesteps of online exploration.

VII. DISCUSSION

We present Learned Efficient Grasp Sets, an algorithm which efficiently explores large sets of grasps by adaptively constructing a small active set of promising grasps. Experiments suggest that LEGS identifies high-performing grasps more efficiently than baseline algorithms across 53 objects in simulation experiments and on three challenging objects in physical trials. We also propose a novel early stopping condition by computing a high-confidence lower bound on the expected grasp performance. Simulation results suggest that this high-confidence lower bound is highly accurate and tight. In future work, we will analyze LEGS to determine how the quality of the Dex-Net prior and the distribution over grasp success probabilities affect its convergence rate. Moreover, we will search for possible ways for LEGS to generalize across different stable poses and objects.

VIII. ACKNOWLEDGMENTS

References

[1] J.-Y. Audibert, S. Bubeck, and R. Munos, "Best arm identification in multi-armed bandits," *COLT 2010 - The 23rd Conference on Learning Theory*, pp. 41–53, Nov. 2010.

[2] D. Berry, R. Chen, A. Zame, D. Heath, and L. Shepp, "Bandit problems with infinitely many arms," *The Annals of Statistics*, vol. 25, Oct. 1997.

[3] D. Brown and S. Niekum, "Efficient probabilistic performance bounds for inverse reinforcement learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 2018.

[4] D. S. Brown, Y. Cui, and S. Niekum, "Risk-aware active inverse reinforcement learning," in *Conference on Robot Learning*, PMLR, 2018, pp. 362–372.

[5] S. Bubeck, T. Wang, and N. Viswanathan, "Multiple identifications in multi-armed bandits," ICML'13, I–258–I–265, 2013.

[6] A. Carpentier and M. Valko, "Simple regret for infinitely many armed bandits," 2015.

[7] D. Chetverikov, D. Svirko, D. Stepanov, and P. Krsek, "The trimmed iterative closest point algorithm," vol. 3, pp. 545–548, 2002.

[8] M. Danielczuk, A. Balakrishna, D. S. Brown, S. Devgon, and K. Goldberg, "Exploratory grasping: Asymptotically optimal algorithms for grasping challenging polyhedral objects," *Conf. on Robot Learning (CoRL)*, 2020.

[9] M. Danielczuk, J. Mahler, C. Correa, and K. Goldberg, "Linear push policies to increase grasp access for robot bin picking," in *2018 IEEE 14th International Conference on Automation Science and Engineering (CASE)*, IEEE, 2018, pp. 1249–1256.

[10] P. Diaconis and D. Ylvisaker, "Conjugate priors for exponential families," *The Annals of statistics*, pp. 269–281, 1979.

[11] C. Eppner and O. Brock, "Visual detection of opportunities to exploit contact in grasping using contextual multi-armed bandits," pp. 273–278, Sep. 2017.

[12] K. Goldberg, B. Mirtich, Y. Zhuang, J. Craig, B. Carlisle, and J. Canny, "Part pose statistics: Estimators and experiments," *IEEE Trans. Robotics and Automation*, vol. 15, no. 5, pp. 849–857, 1999.

[13] R. D. Heide, J. Cheshire, P. M'enard, and A. Carpentier, "Bandits with many optimal arms," *ArXiv*, vol. abs/2103.12452, 2021.

[14] M. Jedor, J. Louëdec, and V. Perchet, "Be greedy in multi-armed bandits," 2021. arXiv: 2101.01086 [cs.LG].

[15] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Scalable deep reinforcement learning for vision-based robotic manipulation," Proceedings of Machine Learning Research, vol. 87, A. Billard, A. Dragan, J. Peters, and J. Morimoto, Eds., pp. 651–673, Oct. 2018.

[16] D. Kalashnkov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," *arXiv*, 2021.

[17] Z. Karnin, T. Koren, and O. Somekh, "Almost optimal exploration in multi-armed bandits," ICML'13, III–1238–III–1246, 2013.

[18] K. Kleeberger, R. Bormann, W. Kraus, and M. F. Huber, "A survey on learning-based robotic grasping," *Current Robotics Reports*, 2020.

[19] M. Laskey, J. Mahler, Z. McCarthy, F. Pokorny, S. Patil, J. van den Berg, D. Kragic, P. Abbeel, and K. Goldberg, "Multi-armed bandit models for 2d grasp planning with uncertainty," pp. 572–579, Aug. 2015.

[20] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015. eprint: https://doi.org/10.1177/0278364914549607.

[21] H. Li, M. Danielczuk, A. Balakrishna, V. Satish, and K. Goldberg, "Accelerating grasp exploration by leveraging learned priors," Nov. 2020.

[22] Q. Lu, M. Merwe, and T. Hermans, "Multi-fingered active grasp learning," pp. 8415–8422, Oct. 2020.

[23] J. Mahler and K. Goldberg, "Learning deep policies for robot bin picking by simulating robust grasping sequences," *Conf. on Robot Learning (CoRL)*, 2017.

[24] 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," *Robotics: Science and Systems (RSS)*, 2017.

[25] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, 2019.

[26] M. Moll and M. A. Erdmann, "Manipulation of pose distributions," *Int. Journal of Robotics Research (IJRR)*, vol. 21, no. 3, pp. 277–292, 2002.

[27] D. Morrison, P. Corke, and J. Leitner, "Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.

[28] D. Morrison, P. Corke, and J. Leitner, "Learning robust, real-time, reactive robotic grasping," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 183–201, 2020.

[29] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," pp. 3406–3413, 2016.

[30] D. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, "A tutorial on thompson sampling," *arXiv preprint arXiv:1707.02038*, 2017.

[31] A. Slivkins, "Introduction to multi-armed bandits," *arXiv preprint arXiv:1904.07272*, 2019.

[32] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.

[33] O. Teytaud, S. Gelly, and M. Sebag, "Anytime many-armed bandits," in *CAP07*, Grenoble, France, 2007.

[34] U. Viereck, A. Pas, K. Saenko, and R. Platt, "Learning a visuo-motor controller for real world robotic grasping using simulated depth images," in *Proceedings of the 1st Annual Conference on Robot Learning*, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., ser. Proceedings of Machine Learning Research, vol. 78, PMLR, 2017, pp. 291–300.

[35] D. Wang, D. Tseng, P. Li, Y. Jiang, M. Guo, M. Danielczuk, J. Mahler, J. Ichnowski, and K. Goldberg, "Adversarial grasp objects," *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, pp. 241–248, 2019.

[36] E. Wang, H. Kurniawati, and D. P. Kroese, "Cemab: A cross-entropy-based method for large-scale multi-armed bandits," M. Wagner, X. Li, and T. Hendtlass, Eds., pp. 353–365, 2017.

[37] Y. Wang, J.-Y. Audibert, and R. Munos, "Algorithms for infinitely many-armed bandits," 2008.

[38] J. Weisz and P. K. Allen, "Pose error robust grasping from contact wrench space metrics," in *2012 IEEE international conference on robotics and automation*, IEEE, 2012, pp. 557–562.

We select the hyperparameters based on the ablation studies done on the *Bar Clamp* object in the Dex-Net 2.0 adversarial objects [24]. Specifically, we perform ablation studies on two hyperparameters: s the strength of GQCNN prior and δ for constructing grasp confidence intervals. For each ablation experiments, we run 10 trials of each algorithm with 10 rollouts per trial, where each trial involves sampling a different reservoir of grasps, and each rollout for a trial involves running a grasp exploration algorithm. All experiments are run over a time horizon of $H = 3000$. Our experiments (Table I) show that this set of hyperparameters tuned on a single object can be applied across many objects.

A. Sensitivity to Prior Strength

The effect of the strength of GQCNN prior is first studied in Li *et al.* [21]. In Danielczuk *et al.* [8], s is set to 5. Our sensitivity experiments show that when more grasps are sampled on each stable pose, $s = 1$ shows the best result. In this set of experiments, $\delta = 0.07$.

Strength of Prior	Bar Clamp
$s = 0$	0.285 ± 0.055
$s = 1$	$\mathbf{0.125 \pm 0.036}$
$s = 2$	0.128 ± 0.038
$s = 3$	0.171 ± 0.041
$s = 4$	0.148 ± 0.036
$s = 5$	0.187 ± 0.042

TABLE II: **Sensitivity to GQCNN prior strength** s: We show the optimality gap (mean $\pm$ standard error).

B. Sensitivity to Confidence Interval Parameter

We also performed sensitivity experiments on δ for constructing confidence intervals for all grasps. Intuitively, using a small δ will lead to a larger confidence interval, which may slow down the updates to the active set. Using a large δ will lead to a smaller confidence interval, which may lead to false positives when identifying grasps with low success rate. In our ablation experiments, we find that $\delta = 0.05$ gives the best performance on the bar clamp object. In this set of experiments, $s = 1$.

Confidence Parameter	Bar Clamp
$\delta = 0.01$	0.129 ± 0.040
$\delta = \mathbf{0.05}$	$\mathbf{0.103 \pm 0.031}$
$\delta = 0.10$	0.156 ± 0.043
$\delta = 0.15$	0.148 ± 0.041
$\delta = 0.20$	0.214 ± 0.051
$\delta = 0.25$	0.174 ± 0.044

TABLE III: **Sensitivity to confidence interval parameter** δ: We show the optimality gap (mean $\pm$ standard error).