

Mechanical Search: Multi-Step Retrieval of a Target Object Occluded by Clutter

Michael Danielczuk^{*1}, Andrey Kurenkov^{*2}, Ashwin Balakrishna¹, Matthew Matl¹, David Wang¹, Roberto Martín-Martín², Animesh Garg², Silvio Savarese², Ken Goldberg¹

Abstract—When operating in unstructured environments such as warehouses, homes, and retail centers, robots are frequently required to interactively search for and retrieve specific objects from cluttered bins, shelves, or tables. *Mechanical Search* describes the class of tasks where the goal is to locate and extract a known target object. In this paper, we formalize Mechanical Search and study a version where distractor objects are heaped over the target object in a bin. The robot uses an RGBD perception system and control policies to iteratively select, parameterize, and perform one of 3 actions – push, suction, grasp – until the target object is extracted, or either a time limit is exceeded, or no high confidence push or grasp is available. We present a study of 5 algorithmic policies for mechanical search, with 15,000 simulated trials and 300 physical trials for heaps ranging from 10 to 20 objects. Results suggest that success can be achieved in this long-horizon task with algorithmic policies in over 95% of instances and that the number of actions required scales approximately linearly with the size of the heap. Code and supplementary material can be found at <http://ai.stanford.edu/mech-search>.

I. INTRODUCTION

In unstructured settings such as warehouses or homes, robotic manipulation tasks are often complicated by the presence of dense clutter that obscures desired objects. Whether a robot is trying to retrieve a can of soda from a stuffed refrigerator or pick a customer’s order from warehouse shelves, the target object is often either not immediately visible or not easily accessible for the robot to grasp. In these situations, the robot must interact with the environment to localize the target object and manipulate the environment to expose and plan grasps (see Figure 1). *Mechanical Search* describes a class of tasks where the goal is to locate and extract the target object, and poses challenges in visual reasoning, task, motion, and grasp planning, and action execution.

Significant progress has been made in recent years on sub-problems relevant to Mechanical Search. Deep-learning methods for segmenting and recognizing objects in images have demonstrated excellent performance in challenging domains [16, 30, 41] and new grasp planning methods have leveraged convolutional neural networks (CNNs) to plan and execute high-quality grasps directly from sensor data [13, 23, 25]. By combining object segmentation and recognition methods with action selectors that can effectively choose between different motion primitives in long horizon sequential tasks, multi-step policies can search for a target object and extract it from clutter.

^{*} Authors have contributed equally and names are in alphabetical order.

¹University of California, Berkeley ²Stanford University

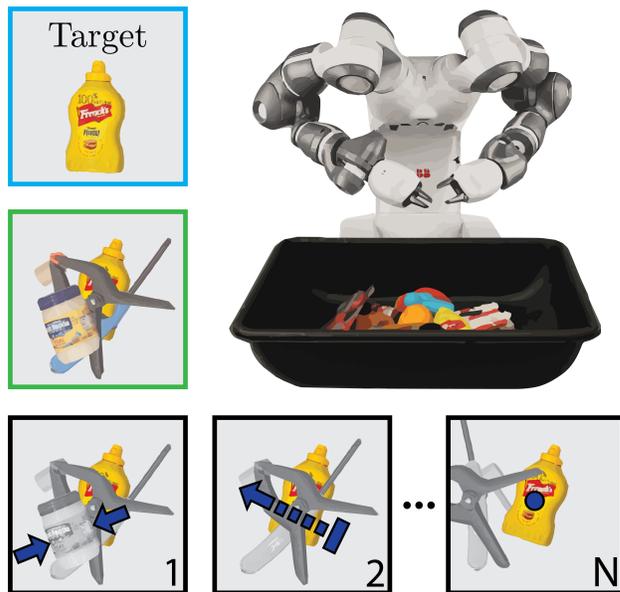


Fig. 1: To locate and extract the target object from the bin, the system selects between 1) grasping objects with a parallel-jaw gripper, 2) pushing objects, or 3) grasping objects with a suction-cup gripper until the target object is extracted, a time limit is exceeded, or no high-confidence push or grasp is available.

In this paper, we propose a framework that integrates perception, action selection, and manipulation policies to address a version of the Mechanical Search problem, with 3 contributions:

- 1) A generalized formulation of the family of *Mechanical Search* problems and a specific version for retrieving occluded target objects from cluttered bins using a series of parallel jaw grasps, suction grasps and pushes.
- 2) An implementation of this version, using depth-based object segmentation, single-shot image recognition, low-level grasp and push planners, and five action selection policies.
- 3) Data from simulation and physical experiments evaluating the performance of the five policies and that of a human supervisor. For simulated experiments, each policy was evaluated on a set of 1000 heaps of 10-20 objects sampled from 1600 3D object models; physical experiments used 50 heaps sampled from 75 common household objects.

II. BACKGROUND AND RELATED WORK

Perception for Sequential Interaction: Searching for an object of interest in a static image is a central problem in active vision [29, 38, 45]. There has also been work

on optimizing camera positioning for improving visual recognition (i.e., *active perception* [1, 2]) and embodied interactions to explore (i.e., *interactive perception* [4, 14]). Mechanical Search differs from prior works in interactive perception in that it deals with long grasping sequences.

Recent deep learning based methods achieve remarkable success in segmentation of RGB [36, 39] and depth images [6], as well as in localizing visual templates in uncluttered [22, 47] and cluttered scenes [30, 41]. Furthermore, one-shot learning approaches using Siamese Networks for matching a novel visual template in images [22, 47] can translate well to pattern recognition in clutter [30, 41]. We build on Mask R-CNN [16] by training a variant for depth-image based instance segmentation and leverage a Siamese network for target template matching for localization.

Grasping and Manipulation in Clutter: Past approaches to this problem can be broadly characterized as model-based with geometric knowledge of the environment [3, 31, 42] and model-free with only raw visual input [20, 24, 40]. Recent studies have leveraged CNNs for casting grasping as a supervised learning problem with impressive results [11, 18, 19, 23, 26, 46, 50]. Pushing and singulation can facilitate grasping in cluttered scenes [5, 8, 17]. Techniques for grasping in clutter, either as open-loop prediction or as closed-loop continuous control, have been studied but have not dealt with the multi-step plans that are critical to attain successful grasps on occluded or inaccessible target objects [19, 32]. In contrast, we formulate Mechanical Search as an interactive search problem in significant clutter, necessitating a multi-step process combining grasping and pushing actions.

Sequential Decision Making: Sequential composition of primitives to enable long-term environment interaction has often been approached through hierarchical decomposition of control policies to manage task complexity. The idea of using hierarchical models for complex tasks has been widely explored in both reinforcement learning and robotics [21, 43, 44]. Training such multi-level models can be computationally expensive and has been limited to either simulated or elementary physical tasks [10, 49].

Search Based Methods: Traditional task planning approaches abstract away perception and focus on high-level task plans and low-level state spaces [12, 42]. For instance, in robotic applications, hierarchical methods have been used to learn task planning strategies while abstracting away low-level motion planning [35, 42, 48]. However, high-level planning requires complete domain specification a priori, and complex geometric and free space reasoning make this approach applicable only to uncluttered environments with few objects, such as a tabletop with one or two objects.

A similar problem has been studied in the context of mobility under problem domains of target-driven and semantic visual navigation [15, 33, 51]. These studies look at finding visual targets in unknown environments without maps through sensory pattern matching. The work by Gupta et al. [14] is the closest to the approach considered in this paper. Their work also considers the problem of searching for a specific

object using pushing and grasping actions, but when the objects are arranged in a shelf. We consider significantly more cluttered settings, while also executing temporally extended manipulation policies.

III. MECHANICAL SEARCH: PROBLEM FORMULATION

In Mechanical Search, the objective is to retrieve a specific target object (x^*) from a physical environment (E) containing a variety of objects X within task horizon H while minimizing time. The agent is initially provided with a specification of the target object in the form of images, text description, a 3D model, or other representation(s). We can frame the general problem of Mechanical Search as a Partially Observable Markov Decision Process (POMDP), defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mathcal{Y})$.

- **States (\mathcal{S}).** A bounded environment E at time t containing N objects $\mathbf{s}_t = \{\mathcal{O}_{1,t}, \dots, \mathcal{O}_{N,t}\}$. Each object state $\mathcal{O}_{i,t}$ includes a ground truth triangular mesh defining the object geometry and pose. Each state also contains the pose and joint states of the robot as well as the poses of the sensor(s).
- **Actions (\mathcal{A}).** A fixed set of parameterized motion primitives.
- **Transitions (\mathcal{T}).** Unknown transition probability distribution $P: \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$.
- **Rewards (\mathcal{R}).** Function given by $R_t = R(\mathbf{s}_t, \mathbf{a}_t) \rightarrow \mathbb{R}$ at time t that estimates the change in probability of successfully extracting the target object $x^* \in X$ within task horizon H .
- **Observations (\mathcal{Y}).** Sensor data, such as an RGB-D image, y_t from robot’s sensor(s) at time t (see Fig. 1).

In this paper, we focus on a specific version of Mechanical Search: extracting a target object specified by a set of k RGB images from a heap of objects in a single bin while minimizing the number of actions needed. For this problem, we precisely specify the observations, the action set, and the reward function. All other aspects of the problem formulation are sufficiently captured by the general POMDP formulation above.

- **Observations.** An RGB-D image from an overhead camera.
- **Actions.**
 - **Parallel Jaw Grasping:** A center point $\mathbf{p} = (x, y, z) \in \mathbb{R}^3$ between the jaws, and an angle in the plane of the table $\phi \in \mathbb{S}^1$ representing the grasp axis [26].
 - **Suction Grasping:** A target point $\mathbf{p} = (x, y, z) \in \mathbb{R}^3$ and spherical coordinates $(\phi, \theta) \in \mathbb{S}^2$ representing the axis of approach of the suction cup [27].
 - **Pushing:** A linear motion of the robot end-effector between two points \mathbf{p} and $\mathbf{p}' \in \mathbb{R}^3$.
- **Reward.** Let v_t , derived from y_t , denote the estimated grasp reliability on the target object. An intuitive reward function would be the increase in estimated grasp reliability on the target object:

$$R(s_t, a_t) = v_{t+1} - v_t$$

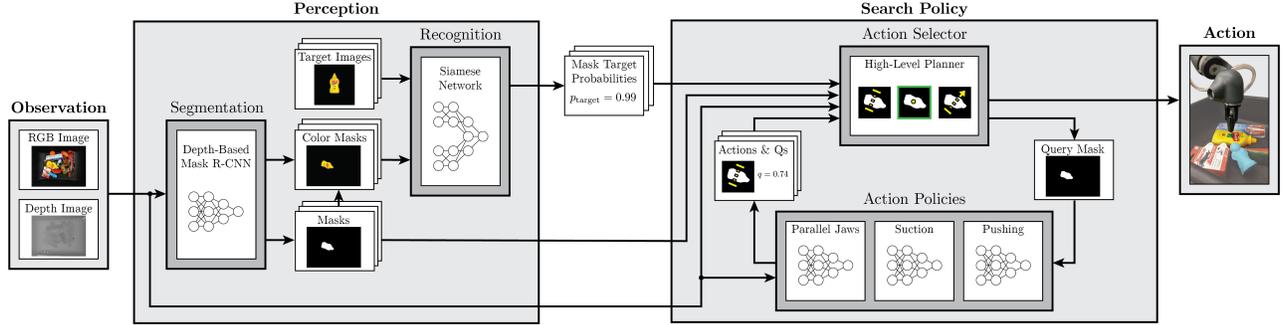


Fig. 2: System architecture. At each timestep, the RGB-D image of the bin is segmented using a variant of Mask R-CNN trained on synthetic depth images. The colorized masks are each assigned a probability of belonging to the target object using a Siamese Network as a pattern comparator. These masks are then fed to an action selector, which chooses which object to manipulate and passes that object mask as context to all the action policies. These policies each compute an action with an associated quality score and pass them back to the action selector, which then chooses an action and executes it on the physical system. This process continues until the target object is retrieved. In simulation, the perception pipeline is removed and planners operate on full-state information rather than object masks.

The policies used in this paper do not directly optimize this reward function because it is difficult to compute; instead, they continue to remove and push objects via heuristic methods until the target object is extracted. In future work, we will develop methods to approximate this function.

IV. PERCEPTION AND DECISION SYSTEM

As shown in Figure 2, we implement the system both in simulation and for physical experiments via a pipeline for perception and policy execution.

A. Perception

The system first processes the RGB-D image into a set of segmentation masks using an object instance segmentation pipeline trained on synthetic depth images. Then, a Siamese network is used to attempt to identify one of the masks as the target object, and a target mask is returned if a high confidence match is found. If no high confidence match is found, the perception system reports that no masks match the target object.

Object Instance Segmentation: We first compute a mask for each object instance. Each mask is a binary image with the same dimensions as the input RGB-D image. These masks are computed with SD Mask R-CNN, a variant of Mask R-CNN trained exclusively on synthetic depth images [9]. It converts a depth image into a list of unclassified binary object masks, and generalizes well to arbitrary objects without retraining. Recent results suggest that depth cues alone may be sufficient for high-performance segmentation, and this network’s generalization capabilities are beneficial in a scenario where only the target object is known and many unknown objects may be present.

Target Recognition: Next, the set of masks is combined with the RGB image to create color masks of each object. Each of the m color masks is cropped, scaled, rotated, and compared to each of the k images in the target object image set using a Siamese network [22]. For each pair of inputs, the Siamese network outputs a recognition confidence value between 0 and 1, with a mask’s recognition confidence score set to the maximum recognition confidence value over the

k target object images. If the mask with the highest score has a score above recognition confidence threshold t_r , the mask is labeled as the target object. Otherwise, we report that no masks match the target object. See the supplementary materials on website for training and implementation details.

B. Search Policy

Given the RGB-D image and the output of the perception pipeline, the system executes the next action in the search procedure by selecting the object to act on and the action to perform on it. Our approach to the version of Mechanical Search described in Section III for bin picking includes searching for actions in three continuous spaces (parallel jaw grasp, suction grasp and push). However, more complex versions of Mechanical Search (e.g., search for an object in a house) could have even more complex search spaces (e.g., navigation). To allow our method to scale to these more complex versions, we propose a hierarchical approach: (1) an action selector that queries a set of action policies on a specific object for a particular action and associated quality metric and (2) action policies that correspond to the possible actions in the problem formulation.

Action Selection: The search policy first determines which object masks to send to the action policies. Then, using the actions and associated quality metric returned by the low level policies, the high level planner determines whether to execute the action in the environment.

The action selector takes as input from the perception system the set of all m visible object masks ($[o_1, \dots, o_m]$), possibly including an object mask that is positively identified as the target object (o_T), from the perception system. It then selects an action policy and a goal object, o_{goal} , from $[o_1, \dots, o_m]$ and sends the action policy a query $q(o_{goal})$. The action policy p_i responds with an action $a_i = p_i(o_{goal})$ and a quality metric $Q(a_i, o_{goal})$ for the action, which is used to decide whether to execute the action.

Action Policies: Each action policy p_i takes as input an object mask from the action selector (o_{goal}) and the RGBD image observation and returns an action $a_i = p_i(o_{goal})$ and a quality metric $Q(a_i, o_{goal})$. In simulation, the object masks and depth images are generated from ground-truth renderings of

each object, while in physical experiments, depth images are obtained using a depth sensor and object masks are generated by the perception pipeline. The set of action policies in our system are:

Parallel Jaw Grasping: In simulation, pre-computed grasps are indexed from a Dex-Net 1.0 parallel-jaw grasping policy [28], and the grasp with the highest predicted quality on o_{goal} is returned as the action along with an associated quality metric. For physical experiments, parallel-jaw grasps are planned using a Dex-Net 2.0 Grasp Quality CNN (GQ-CNN) [25]. To plan grasps for a single object in a depth image, grasp candidate sampling is constrained to the goal object’s segmentation mask. The GQ-CNN evaluates each candidate grasp and returns the grasp with the highest predicted quality and its associated quality metric.

Suction Grasping: For simulation experiments, grasp planning is done with a Dex-Net 1.0 suction grasping policy [28]. For physical experiments, suction cup grasps are planned with a Dex-Net 3.0 GQ-CNN [27], with mask-based constraints to plan grasps only on the goal object’s segmentation mask. The GQ-CNN evaluates each candidate grasp and returns the grasp with the highest predicted quality and its associated quality metric.

Pushing: The pushing action policy, similar to that in [8], selects \mathbf{p}' as the most free point in the bin. This point is computed by taking the signed distance transform of a binary mask of the bin walls and objects, finding the pixel with the maximum signed distance value, and deprojecting that pixel back into \mathbb{R}^3 . Given an object to push, \mathbf{p} is then selected so that the gripper is not in collision at \mathbf{p} , the line from \mathbf{p} to \mathbf{p}' passes through the object’s center of mass, and the push direction is as close as possible to the direction of the most free point in the bin. The pushing policy returns the push satisfying the above constraints as its action if one exists. The returned quality metric is 1 if a valid push exists and 0 if not.

V. ACTION SELECTION POLICIES

All action selection methods use input from the perception system to generate a specific object priority list. Each action selection method generates a priority list in a different way but all have the same action execution criteria. For all action selection methods described here, a grasp action is executed if the quality metric returned by the action policy exceeds $t(o)$, the grasp confidence threshold for object mask o . The grasp confidence threshold for the object mask positively identified as the target object o_T is given by $t(o_T) = t_{\text{thresh}}$. For policies without pushing, $t(o) = t_{\text{thresh}}, \forall o$, while for policies with pushing, $t(o) = t_{\text{high}}, \forall o \neq o_T$. Policies with pushing can be more conservative in their choice of grasps, so they use a higher grasp confidence threshold t_{high} for non-target objects. A push action is performed if a valid push is found (quality 1). Details on parameters used can be found in the supplementary material.

Each action selection method iterates through its priority list, queries the grasping action policies for each object mask, and executes the returned action with the highest quality

metric among the two grasping policies if it satisfies the action execution criteria. If the target object is grasped, the policy terminates and reports a success. If no grasping action satisfies the criteria and the policy does not have pushing, the policy terminates and reports a failure. If the policy does have pushing, it iterates through its priority list, queries the pushing action policy for each object mask, and executes the first action that satisfies the criteria. If no pushing action satisfies the criteria, or if a pushing action has been selected more than three consecutive times, the policy terminates with a failure.

Action Selection Methods: The action selection methods are distinguished by whether or not they have pushing as an available action policy and by their generated object priority list:

- 1) **Random Search:** Prioritizes objects randomly, with no preference for the target object mask (o_T).
- 2) **Preempted Random Search (with and without pushing):** Always prioritizes o_T and prioritizes other objects randomly.
- 3) **Largest-First Search (with and without pushing):** Always prioritizes o_T and ranks the other objects by their visible area. If the target object isn’t visible, this strategy will increase the likelihood of removing objects that may be occluding the target object.

Termination Criteria: In addition to the termination criteria outlined above (terminate and return success if target object grasped, return failure if no good grasp/valid push found), we impose two more termination conditions on our policies which cause them to return a failure: (1) $2N$ timesteps have elapsed, where N is the initial number of objects in the bin and (2) The target object is inadvertently removed from the work space when another object is grasped or pushed.

VI. EXPERIMENTS

A. Simulation

Heap Generation: Three datasets of simulated heaps are generated, each containing 1000 heaps of N objects, for $N \in \{10, 15, 20\}$. Then, using the Bullet Physics Engine [7], sampled objects are dropped one by one into the bin, and the target object is chosen to be the most occluded object. Please refer to the supplementary material for further details.

Rollouts: To simulate grasp actions, we use the same approach as in [24]: using wrench space analysis, we determine whether or not an object can be lifted from the heap [34, 37]. If the object can be lifted, a constant upward force is applied to the object’s center of mass until it leaves the bin, and the remaining objects are allowed to come to rest. To simulate push actions, we check that the gripper can be placed in the starting location without collisions, and only execute pushes if this is the case. Then, we place a 3D model of the closed gripper in the physics simulator and move it from the start point to the end point of the push, as in [8].

B. Physical

Heap Generation: We randomly sample 50 heaps of 15 items each from a set of 75 common household objects with

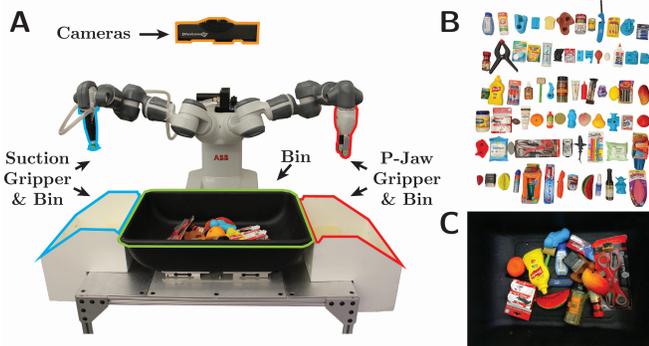


Fig. 3: (A) Front view of the robot and bin setup. The black bin is the primary bin in which heaps are initialized, and the white bins provide space for the robot to deposit grasped items. (B) The 75 objects used in physical experiments. (C) A sample heap of 15 objects used in the physical experiments.

relatively simple shapes, such as boxes and cylinders, as well as more complex geometries, such as plastic climbing holds and scissors (see Figure 3). We also include several 3D-printed items, which present a challenge for both segmentation and target object recognition due to their unusual shapes and uniform texture. A target object is chosen at random from each 15 item heap. Then, in order to generate adversarial bin configurations, each rollout is initialized by first shaking the target object in a box to randomize its pose and dumping into the center of the bin, and then shaking the other fourteen objects and pouring them over the target object.

Policy Rollouts: We execute pushing and grasping actions on an ABB YuMi robot equipped with suction-cup and parallel-jaw grippers (see Figure 3). Actions generated by the search policy are transformed into a sequence of poses for the robot’s end-effectors, and we use ABB’s RAPID linear motion planner and controller to execute these motions.

Human Supervisor Rollouts: For comparison, we also benchmark a human supervisor’s performance as an action selector. At each timestep, the human is asked to draw a mask in the scene on which to plan a push or a grasp. Then, grasps and pushes are planned and executed on the specified mask with the same action primitives described above (parallel jaw grasps, suction grasps, linear pushes). Thus, the human is limited by the available action primitives, but is allowed to use their own judgement for perceptual reasoning and high level action planning.

C. Evaluation Metrics

We evaluate each policy according to its reliability and efficiency in target object retrieval. Reliability is defined as the frequency at which the target object is successfully extracted, while efficiency is defined as the mean number of actions taken to successfully extract the target object. For each experiment, we recorded the number of successes and failures, as well as statistics regarding the number of actions taken.

VII. RESULTS

A. Simulation Results

We tested each action selection policy in simulation with heaps generated using the method described in VI and running

until termination on each heap. A total of 15,000 simulation experiments were conducted over all policies. The results for each of the five policies are shown in Figure 4, and a detailed breakdown of each policy can be found in the supplementary material. Figure 4(A) shows the mean number of actions needed for each policy as a function of heap size. These results suggest that extracting the target object becomes more difficult as heap size increases and the mean number of actions needed for each policy appears to scale linearly with heap size, although the rate of increase is not constant across policies.

Figure 4(B) shows cumulative successful extractions for a given amount of actions (number of successful extractions in that many actions or fewer). All policies have success rates of 90% or higher on 15 object heaps. The cumulative success plot suggests that grasping the target object when possible provides improvement over the random policy, and that prioritizing larger object masks when the target object is inaccessible further increases efficiency. The largest-first policies successfully extract the target object within 5 or fewer actions on 50% of the heaps, while the preempted random and random policies only do so for 30% and 10% of heaps respectively.

Results also suggest that pushing can increase overall success rate, as policies that included pushing succeeded on at least 3% more heaps than those without pushing. For policies with pushing, only 5% of all actions attempted are pushes, as opposed to 29% parallel jaw grasps and 66% suction grasps, on ten object heaps. The percentage of push actions decreases further when increasing the number of objects to just 3% for 15 object heaps. We suspect that the reason pushes are selected so rarely is because the pushing primitive is designed to execute a sequence of linear pushes to singulate a particular object, rather than flatten a heap of objects. While the former directly addresses the objective of the push, the latter is much easier to achieve in practice, since in many cases, especially with many objects in the bin, it may be almost impossible to plan collision-free pushes that successfully singulate the object of interest.

In simulation, failures account for 6-12% of all rollouts. These failures fall into three categories: 1) the policy fails to plan an action (e.g., no actions are available on any remaining objects with a quality metric above the threshold), 2) the target object is inadvertently removed from the bin, despite an attempted action that was not a grasp on the target object, or 3) the policy reaches the maximum number of timesteps given to extract the object ($2N$ timesteps, where N is the initial number of objects in the heap).

Failure mode (1) accounts for 85-90% of all failures, while (2) accounts for nearly all of the remaining failures. Mode (1) typically happens because objects are moved to the corners of the bin, making it difficult to plan collision-free grasps. Mode (2) occurs more often for policies that include pushing actions, since pushing in clutter can occasionally lead to objects being pushed up and over the bin walls. It is also possible that the target object is removed from the work space when another object is grasped. Mode (3) never actually occurs in any

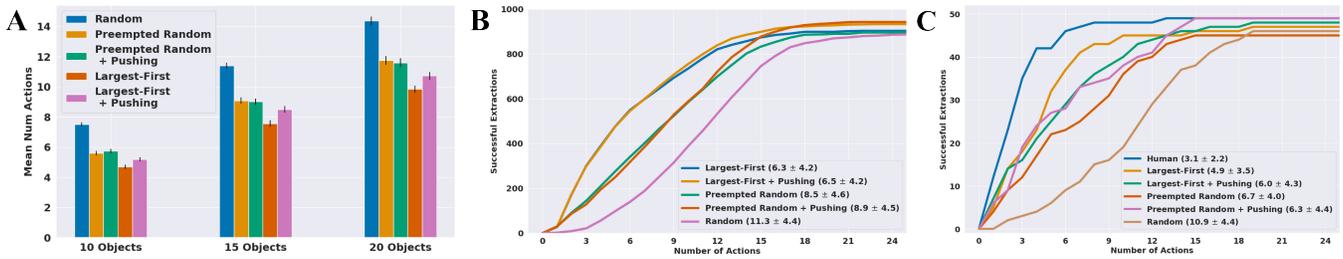


Fig. 4: Performance of policies on (A) simulated heaps of 10, 15, and 20 objects over a total of 15,000 simulated rollouts and 300 physical rollouts, (B) simulated heaps of 15 objects, and (C) real heaps of 15 objects. The largest-first search policies are the most efficient, and are able to extract the target object in the least number of actions. All policies have similar reliability, although pushing shows potential to avoid more failures in simulation. The human was allowed to look at the RGBD image inputs and choose an object to push or grasp. Means and standard deviations for successful extractions are shown in parentheses for each policy.

of our experiments, and the maximum timesteps cutoff is intentionally set high to exhaust the policy of actions.

B. Physical Results

Figure 4(C) shows results on the physical system. A total of 300 physical experiments were conducted over all policies. All policies retrieved the target object within the given number of timesteps at least 90% of the time, and success rates were not statistically different between policies. However, the cumulative success curves suggest similar trends to those seen in the simulation results, with the largest-first policies outperforming the preempted random and random policies in terms of efficiency. The largest-first policies again successfully extract the target object within 5 or fewer actions on 50% of the heaps, while the preempted random and random policies only do so for 40% and 10% of heaps respectively.

Stochasticity in the initial bin state in physical experiments can result in varying difficulty for different policies on the same heap. For example, a target object may be completely covered by other objects when one policy is presented with a given heap, but for another policy, the target may be partially or fully visible in the initial state. Thus, policies may occasionally get “lucky” or “unlucky” with respect to the target object visibility in the initial state, which may account for some increased variance in the physical results.

Failure cases for the physical heaps are very similar to those in simulation: 93% of failures arise from the policy being unable to plan an action. However, in physical experiments, out of the 300 heaps evaluated for all policies, only 1 rollout failed due to timing out. Failure to plan actions is almost always due to the target object lying flat on the bottom of the bin (e.g., the dice, sharpie pens, or another blister-pack object), making it difficult to obtain accurate segmentation. Another common reason for failure to plan actions is when no mask is identified as the target object, which often occurs for 3D printed objects.

C. Action-Limited Human Supervisor

The human supervisor outperforms all policies presented here, requiring an average of just 3.1 actions to extract the target object due to more intelligent action selection. Specifically, we noticed that a human operator chose to push far more frequently (26% of all actions, compared to 6% for the other action policies with pushing), especially when

objects were heaped in the center of the bin and the target was not visible. These pushes tend to spread many objects out over the bottom of the bin, as opposed to a grasping action that would remove only a single object from the top of the heap.

VIII. DISCUSSION AND FUTURE WORK

We present a general formulation for mechanical search problems and describe a framework for solving the specific problem of extracting a target object from a cluttered bin. While the best action selection method (largest-first) is much more efficient than random search and provides a solid baseline, a human selecting the low-level actions can still achieve 37% higher efficiency by pushing significantly more effectively and often. We will explore how reinforcement learning in simulation can address this gap.

In future work, the action primitives used (grasp, suction, push) can also be improved. We conjecture that more effective push primitives can be learned from simulation.

IX. ACKNOWLEDGMENTS

This work is partially supported by a Google Focused Research Award and was performed jointly at the AUTOLAB at UC Berkeley and at the Stanford Vision & Learning Lab, in affiliation with the Berkeley AI Research (BAIR) Lab, Berkeley Deep Drive (BDD), the Real-Time Intelligent Secure Execution (RISE) Lab, and the CITRIS “People and Robots” (CPAR) Initiative. Authors were also supported by the SAIL-Toyota Research initiative, the Scalable Collaborative Human-Robot Learning (SCHoOL) Project, the NSF National Robotics Initiative Award 1734633, and in part by donations from Siemens, Google, Amazon Robotics, Toyota Research Institute, Autodesk, ABB, Knapp, Loccioni, Honda, Intel, Comcast, Cisco, Hewlett-Packard and by equipment grants from PhotoNeo, and Nvidia. This article solely reflects the opinions and conclusions of its authors and do not reflect the views of the Sponsors or their associated entities. We thank our colleagues who provided helpful feedback, code, and suggestions, in particular Jeff Mahler.

REFERENCES

- [1] A. Aydemir, K. Sjö, J. Folkesson, A. Pronobis, and P. Jensfelt, "Search in the real world: Active visual object search based on spatial relations", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2011, pp. 2818–2824.
- [2] R. Bajcsy, "Active perception", *Proceedings of the IEEE*, vol. 76, no. 8, pp. 966–1005, 1988.
- [3] D. Berenson and S. S. Srinivasa, "Grasp synthesis in cluttered environments for dexterous hands", in *Humanoid Robots, 2008. Humanoids 2008. 8th IEEE-RAS International Conference on*, IEEE, 2008, pp. 189–196.
- [4] J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, and G. S. Sukhatme, "Interactive perception: Leveraging action in perception and perception in action", *IEEE Trans. Robotics*, vol. 33, no. 6, pp. 1273–1291, 2017.
- [5] L. Chang, J. R. Smith, and D. Fox, "Interactive singulation of objects from a pile", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2012, pp. 3875–3882.
- [6] X. Chen, K. Kundu, Y. Zhu, H. Ma, S. Fidler, and R. Urtasun, "3d object proposals using stereo imagery for accurate object class detection", *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 5, pp. 1259–1272, 2018.
- [7] E. Coumans and Y. Bai, *Pybullet, a python module for physics simulation, games, robotics and machine learning*, <http://pybullet.org/>, 2016–2017.
- [8] M. Danielczuk, J. Mahler, C. Correa, and K. Goldberg, "Linear push policies to increase grasp access for robot bin picking", in *Proc. IEEE Conf. on Automation Science and Engineering (CASE)*, IEEE, 2018.
- [9] M. Danielczuk, M. Matl, S. Gupta, A. Li, A. Lee, J. Mahler, and K. Goldberg, "Segmenting unknown 3d objects from real depth images using mask r-cnn trained on synthetic point clouds", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2019.
- [10] Y. Duan, M. Andrychowicz, B. C. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-Shot Imitation Learning", *arXiv preprint arXiv:1703.07326*, 2017.
- [11] K. Fang, Y. Zhu, A. Garg, A. Kurenkov, V. Mehta, L. Fei-Fei, and S. Savarese, "Learning task-oriented grasping for tool manipulation from simulated self-supervision", *arXiv preprint arXiv:1806.09266*, 2018.
- [12] R. E. Fikes and N. J. Nilsson, "Strips: A new approach to the application of theorem proving to problem solving", *Artificial intelligence*, vol. 2, no. 3-4, pp. 189–208, 1971.
- [13] M. Gualtieri, A. ten Pas, K. Saenko, and R. Platt, "High precision grasp pose detection in dense clutter", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, IEEE, 2016, pp. 598–605.
- [14] M. Gupta, T. Rühr, M. Beetz, and G. S. Sukhatme, "Interactive environment exploration in clutter", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, IEEE, 2013, pp. 5265–5272.
- [15] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation", *arXiv preprint arXiv:1702.03920*, vol. 3, 2017.
- [16] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn. arxiv preprint arxiv: 170306870", 2017.
- [17] T. Hermans, J. M. Rehg, and A. Bobick, "Guided pushing for object singulation", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, IEEE, 2012, pp. 4783–4790.
- [18] E. Jang, S. Vijayanarasimhan, P. Pastor, J. Ibarz, and S. Levine, "End-to-end learning of semantic grasping", in *Conf. on Robot Learning (CoRL)*, 2017.
- [19] 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.
- [20] D. Katz, A. Venkatraman, M. Kazemi, J. A. Bagnell, and A. Stentz, "Perceiving, learning, and exploiting object affordances for autonomous pile manipulation", *Autonomous Robots*, vol. 37, no. 4, pp. 369–382, 2014.
- [21] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey", *Int. Journal of Robotics Research (IJRR)*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [22] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition", in *ICML Deep Learning Workshop*, vol. 2, 2015.
- [23] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps", *Int. Journal of Robotics Research (IJRR)*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [24] J. Mahler and K. Goldberg, "Learning deep policies for robot bin picking by simulating robust grasping sequences", in *Conf. on Robot Learning (CoRL)*, 2017, pp. 515–524.
- [25] 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", *Proc. Robotics: Science and Systems (RSS)*, 2017.
- [26] —, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics", *CoRR*, vol. abs/1703.09312, 2017. arXiv: 1703.09312.
- [27] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dex-net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2018, pp. 1–8.
- [28] 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 *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2016, pp. 1957–1964.
- [29] C. Michaelis, M. Bethge, and A. Ecker, "One-shot segmentation in clutter", in *Proc. Int. Conf. on Machine Learning*, J. Dy and A. Krause, Eds., ser. Proceedings of Machine Learning Research, vol. 80, Stockholm, Sweden: PMLR, 2018, pp. 3549–3558.
- [30] C. Michaelis, M. Bethge, and A. S. Ecker, "One-shot segmentation in clutter", *arXiv preprint arXiv:1803.09597*, 2018.
- [31] M. Moll, L. Kavraki, J. Rosell, et al., "Randomized physics-based motion planning for grasping in cluttered and uncertain environments", *IEEE Robotics & Automation Letters*, vol. 3, no. 2, pp. 712–719, 2018.
- [32] D. Morrison, P. Corke, and J. Leitner, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach", *arXiv preprint arXiv:1804.05172*, 2018.
- [33] A. Mousavian, A. Toshev, M. Fisher, J. Kosecka, and J. Davidson, "Visual representations for semantic target driven navigation", *arXiv preprint arXiv:1805.06066*, 2018.
- [34] A. ten Pas and R. Platt, "Using geometry to detect grasp poses in 3d point clouds", in *Robotics Research*, Springer, 2018, pp. 307–324.
- [35] C. Paxton, F. Jonathan, M. Kobilarov, and G. D. Hager, "Do what i want, not what i did: Imitation of skills by planning sequences of actions", in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
- [36] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár, "Learning to refine object segments", in *European Conference on Computer Vision*, Springer, 2016, pp. 75–91.
- [37] D. Prattichizzo and J. C. Trinkle, "Grasping", in *Springer handbook of robotics*, Springer, 2008, pp. 671–700.
- [38] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection", in *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.
- [39] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks", in *Proc. Advances in Neural Information Processing Systems*, 2015, pp. 91–99.
- [40] A. Saxena, L. L. Wong, and A. Y. Ng, "Learning grasp strategies with partial shape information.", in *AAAI*, vol. 3, 2008, pp. 1491–1494.
- [41] A. Shaban, S. Bansal, Z. Liu, I. Essa, and B. Boots, "One-shot learning for semantic segmentation", *arXiv preprint arXiv:1709.03410*, 2017.
- [42] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel, "Combined task and motion planning through an extensible planner-independent interface layer", in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, IEEE, 2014, pp. 639–646.
- [43] J. Sung, B. Selman, and A. Saxena, "Learning sequences of controllers for complex manipulation tasks", in *Proc. Int. Conf. on Machine Learning*, Citeseer, 2013.
- [44] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning", *Artificial intelligence*, vol. 112, no. 1-2, pp. 181–211, 1999.
- [45] K. E. Van de Sande, J. R. Uijlings, T. Gevers, and A. W. Smeulders, "Segmentation as selective search for object recognition", in *Proc.*

- IEEE Int. Conf. on Computer Vision (ICCV)*, IEEE, 2011, pp. 1879–1886.
- [46] U. Viereck, A. t. Pas, K. Saenko, and R. Platt, “Learning a visuomotor controller for real world robotic grasping using simulated depth images”, *arXiv preprint arXiv:1706.04652*, 2017.
- [47] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra, *et al.*, “Matching networks for one shot learning”, in *Proc. Advances in Neural Information Processing Systems*, 2016, pp. 3630–3638.
- [48] J. Wolfe, B. Marthi, and S. J. Russell, “Combined task and motion planning for mobile manipulation”, in *ICAPS*, 2010.
- [49] D. Xu, S. Nair, Y. Zhu, J. Gao, A. Garg, L. Fei-Fei, and S. Savarese, “Neural task programming: Learning to generalize across hierarchical tasks”, in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2018.
- [50] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, *et al.*, “Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching”, *arXiv preprint arXiv:1710.01330*, 2017.
- [51] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, “Target-driven visual navigation in indoor scenes using deep reinforcement learning”, in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 3357–3364.