

Interactive Perception in Clutter

Megha Gupta
Computer Science Department
University of Southern California
Los Angeles, California 90089
Email: meghagup@usc.edu

Gaurav S. Sukhatme
Computer Science Department
University of Southern California
Los Angeles, California 90089
Email: gaurav@usc.edu

Abstract—In robotic manipulation tasks, perception traditionally precedes the act of manipulation, after which the two are typically interleaved. Perception thus enables efficient and effective manipulation. This paper explores the idea that manipulation enables efficient and effective perception, which in turn improves performance on the overall robotic manipulation task. We study this idea in cluttered environments in two contexts - sorting small objects on a crowded tabletop and exploration of small occluded environments (e.g., the shelf of a home refrigerator). For object sorting, we present a robust pipeline that uses two simple motion primitives to manipulate a cluttered scene of Duplo bricks in ways that improve the robot’s perceptual, (and hence) grasping and sorting abilities. Using the PR2, we experimentally compare brick sorting with and without manipulation aids explicitly introduced to improve perception. We show substantial improvement on the overall task (sorting) by inserting manipulation aids into the task operation. For exploring a cluttered environment, we describe a strategy that uses adaptive look-ahead planning interleaved with object manipulation to improve perception. We present preliminary simulation results using this planner.

I. INTRODUCTION

Picture this. Your 8-year old kid comes to you asking for his jigsaw puzzle. You don’t exactly remember where you have kept it and you go to the toy cupboard looking for it. It has been a while since you last saw it and you have to displace and remove many things in the cupboard before you finally find it. But now your kid wants you to help him solve the puzzle too. You both decide that sorting the pieces first is a good idea - finding the edge pieces and sorting them roughly by color is probably a good starting point.

As it turns out, we need to search for and sort objects more often than we realize. Think about looking for your car keys, that book that you need to return to the library, a matching pair of socks in your drawer, sorting clothes for laundry, your grocery, your snail mail - the examples are endless. These tasks are made tougher by the fact that our homes are cluttered - even those that are very organized may have a lot of objects. The good news is that we usually succeed at them in reasonable amounts of time. One of the common factors that facilitates or even enables both searching and sorting is our ability to manipulate the world around us.

One of the key challenges of household robotics is to make robots self-sufficient and capable of understanding their surroundings well. An important aspect of this self-sufficiency is the ability to deal with clutter. Traditionally, perception guides manipulation in robotic manipulation tasks. Having

sophisticated perceptual abilities is, thus, very important to deal with the large variety and complexity of objects that a robot may encounter. However, clutter introduces uncertainty and errors in robotic perception leading to failures in a manipulation task.

Interactive perception, or using manipulation of the world to enable efficient perception, could serve as an elegant tool in such scenarios. Effective perception would, in turn, improve the performance of the overall robotic manipulation task. Humans use it all the time - in our jigsaw puzzle example, you had to displace objects in the cupboard to search for another object, and then had to shuffle through the pieces of the jigsaw puzzle to sort them. An appealing trait of these manipulation maneuvers (other than their effectiveness) is that they are few and simple e.g., picking up an object, pushing things aside, and shuffling through them. This simplicity makes it natural to ask - can a robot do the same in clutter?

It is this idea of interactive perception in clutter that we address in this paper. Being cluttered, the environment is only partially observable. The robot can manipulate the world to observe it more and better, and re-plan based on the new observations. This setup suggests an observe-plan-act strategy where perception, planning, and action have to be interleaved.

We test such an approach in two different contexts introduced earlier - searching for an object and sorting objects in clutter. In particular, we present algorithms for sorting small objects (Duplo bricks) on a crowded tabletop by color or size; and exploring small occluded environments (e.g., a bookshelf) to search for or count objects. We demonstrate through experiments on the PR2 that better Duplo sorting results are obtained when cluttered parts of the scene are first ‘manipulated’ using simple actions, as opposed to when bricks are attempted to be picked up directly from clutter. For environment exploration, we propose an adaptive look-ahead exploration algorithm that suggests which objects in the world should be displaced and where. We present preliminary simulation results for the same.

II. RELATED WORK

Sequential decision making can be broadly classified into two types - one where the state is fully observable and the other, where it is partially observable or unknown. The first class of problems can be solved by Markov Decision Processes or MDPs. When the actions are deterministic, we get

a deterministic decision tree, and finally when state transitions are deterministic as well, offline planning algorithms like Breadth First Search and Depth First Search can be employed.

However, when the states and state transitions are uncertain, a popular planning framework is POMDPs (Partially Observable Markov Decision Processes) that maintain a distribution over belief states as opposed to states [18]. However, a POMDP becomes intractable very quickly as the size of the state space grows. Heuristic-based algorithms like A^* , D^* , and D^* -lite ([14], [25]) work well when the goal state is known and a good heuristic is used.

The Duplo sorting problem has fully observable states with deterministic actions while the environment exploration problem has partially observable states with multiple actions. Both problems, however, have uncertain state transitions and therefore, require perception and planning interleaved with manipulation.

Several efforts have been directed at robust robot operation in cluttered environments. Most of the work in this area either focuses on successful manipulation in clutter ([4], [11], [5]) or on accurate perception in clutter ([24], [9], [7], [15]). There has also been research on building accurate object models from perceptual data so that grasp planning can benefit from it. These approaches study perception techniques to aid manipulation while we are aiming at the converse.

Traditionally, research on mobile manipulation considers collisions with obstacles unacceptable. More recently, there has been some work on *interactive perception* ([6], [17], [23], [13], [3]) where the robot actively manipulates the world to complete the overall task at hand. Our work in this paper falls in this category.

Parts sorting is an old field of research and highly accurate vibration-based parts feeders have been designed in industry ([2], [12]). However, these methods assume availability of very specific equipment or require parts to be fed to the feeder one by one. Our method does not have these requirements.

The problem of exploration and object search in different kinds of environments has also been studied extensively. Algorithms have been designed in the context of *active visual search* ([1], [26], [19]), *pursuit-evasion* ([20], [8]), and *search and rescue operations* ([16]). Our problem of environment exploration requires planning under uncertainty but the goal state is unknown, and there is no prior on object locations. We, therefore, propose adaptive look-ahead planning interleaved with manipulation of objects in the environment for exploration.

III. DUPLO BRICK SORTING IN CLUTTER

A. Algorithm

All our Duplo sorting experiments were done on the PR2 with a head-mounted Kinect. Fig. 1 shows our experimental setup. Duplos are scattered on a table and the PR2 is positioned at one of the edges of the table with its head looking down. Bins for collecting sorted bricks are placed on one side of the table at known locations. For the rest of the paper, we will assume the tabletop to lie in the XY -plane with the X

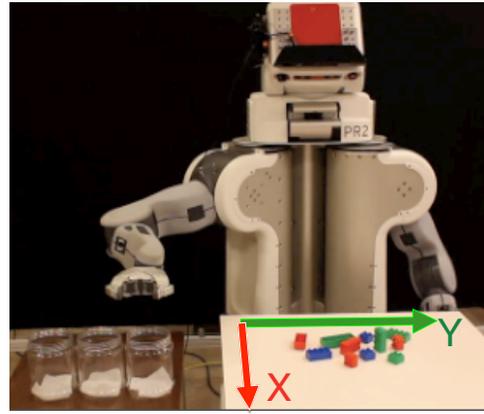


Fig. 1: Experimental setup. The PR2 with a head-mounted Kinect views Duplo bricks on a flat table-top. The task is to sort the bricks by color and size into the three bins placed next to the table.

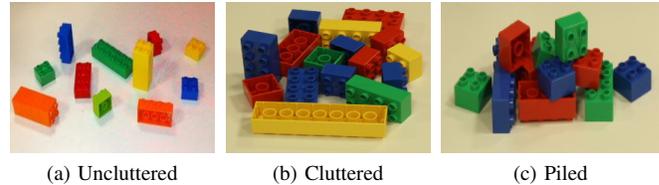


Fig. 2: The three different states a region could be in.

and Y axes as shown in Fig. 1. For this setup, we combine perception and manipulation to sort the Duplos by color or size (length) in separate bins and eventually, clear the table. Algorithmically, our pipeline is simple:

- 1) Divide the observed tabletop into non-overlapping regions with each region containing closely placed Duplos (details in section III-B). For each region, do the following: (a) Assign a state to the region based on the degree of clutter in the region. (b) Manipulate the region using the appropriate manipulation primitive based on its state.
- 2) Repeat till the table is cleared.

The states are defined to be the following:

- *Uncluttered*: No two Duplo bricks in the region are in contact with each other.
- *Cluttered*: There are bricks that are close to or in contact with at least one other brick. However, all of them lie directly on the table.
- *Piled*: The bricks are piled up, i.e., some of them lie on top of other bricks.

Each state has an associated action:

- *Uncluttered*: Pick and place every brick in the corresponding bin and clear the region.
- *Cluttered*: Spread out the bricks (increase their average separation) to make the region uncluttered.
- *Piled up*: Decompose the pile.

Implementation details of the algorithm follow in the next section.

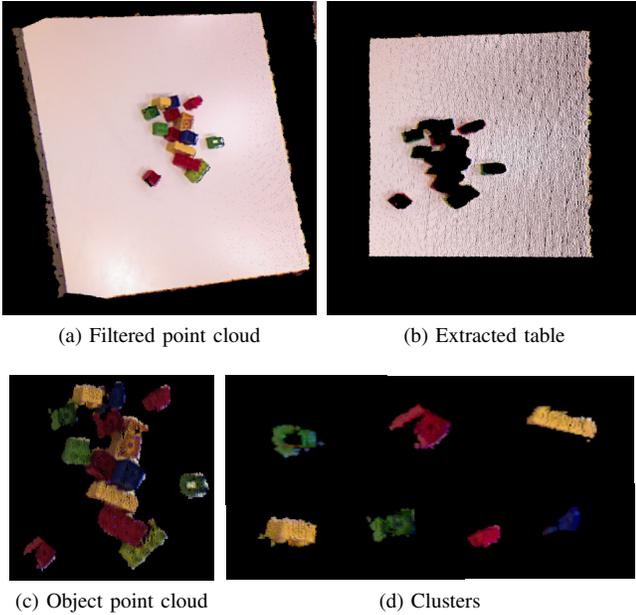


Fig. 3: Point cloud extraction for every Duplo from the point cloud received from Kinect.

B. Implementation: Perception

The perception pipeline receives a point cloud from the head-mounted Kinect sensor and processes it as follows:

1) *Object Cloud Extraction*: Standard algorithms available in the Point Cloud Library (PCL) [22] are used for preprocessing. Planar segmentation of the filtered point cloud is used to separate the tabletop and the objects from the point cloud. The point cloud for each Duplo brick is then obtained from the object cloud using Euclidean color-based clustering algorithm which assigns the same cluster to neighboring points that are also of the same color. Since every Duplo brick is a single-color object, this method allows two bricks of different colors to be distinguished from each other even if they are in contact and their point clouds overlap. Fig. 3 shows the results of these pre-processing steps for a point cloud.

2) *Dividing into Regions*: The object point cloud is processed to obtain different regions such that the Duplos within a region are all closer than a pre-defined threshold to each other while Duplos in different regions are at least as far as the threshold. The Euclidean clustering algorithm available in PCL [21] is used for this. Fig. 4 shows the regions generated by the algorithm for two different scenes. Learning the number of regions from the point cloud makes the algorithm adaptive and allows the best manipulation decisions to be taken irrespective of the distribution of the Duplos on the table.

3) *Assigning States to Regions*: For each region, we check if the Duplos are in contact and what the height of the point cloud is and based on this, the region is classified as *uncluttered*, *cluttered*, or *piled*.

C. Implementation: Manipulation Primitives

1) *Pick and Place*: For an uncluttered region, a slightly modified PR2 tabletop manipulation pipeline, which is avail-

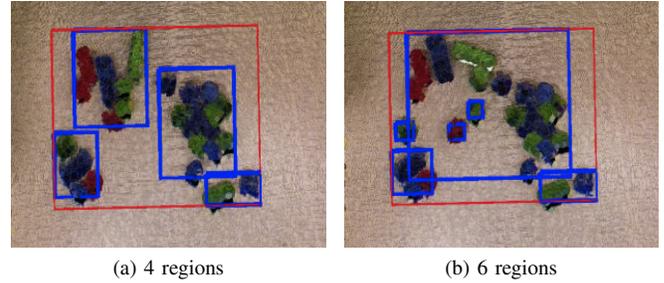


Fig. 4: Snapshots from RViz showing different regions identified by the perception pipeline. Bricks within a region are close to each other while any two bricks in different regions are a minimum distance apart. The red rectangle shows the bounding box for the object point cloud. The blue rectangles show the bounding boxes for the regions.

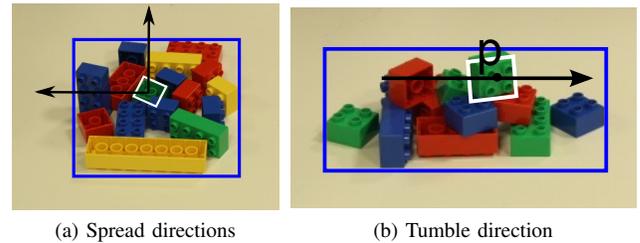


Fig. 5: Calculating spread and tumble directions. The blue rectangles show the bounding boxes for the regions and the white rectangles mark the Duplo that is manipulated. Spread and tumble directions are chosen as the ones shown by black arrows.

able as a ROS package, is used to sort Duplos into bins. Collisions with other bricks and the table are allowed. The robot attempts to pick up every cluster in the region's point cloud. Once a brick is picked up, the arm moves to a pre-defined bin position based on the color or size of the brick and drops it in the bin. This action is likely to make the region empty unless a grasp fails.

2) *Spread*: For a cluttered region, first the brick that is in contact with the maximum number of other bricks is located. Perturbing this brick from its original position moves the neighboring bricks as well and may result in isolating several bricks for easy pick up. Let us denote this brick by B . The robot places its fingers on the center of B and moves it once along the X -axis and then along the Y -axis. The directions (forward/backward along X axis and left/right along Y axis) of these movements are calculated from the point cloud of the region and chosen to be the ones that are more likely to take the brick away from others. Specifically, the directions are such that the extents of the point cloud around B are smaller in those directions. See Fig. 5a for an example. This action is likely to make the region less cluttered than before.

3) *Tumble*: For a piled up region, the centroid of the topmost Duplo in the region is located. Let us call this point as p and the direction that the X axis points in as forward. Three waypoints are defined for the robot arm: a point in front of p , p , and a point behind p . The robot moves its hand across these waypoints in an attempt to tumble the pile. The direction of this movement is the one of smaller extent of the point cloud

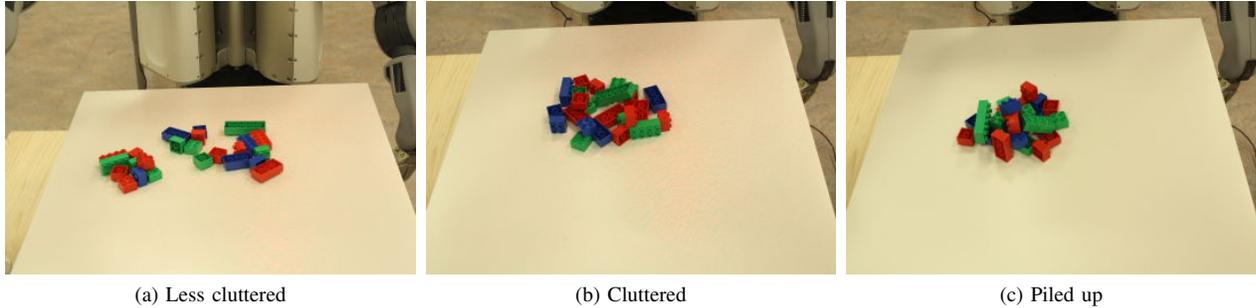


Fig. 6: Scenes of varying degree of clutter to compare sorting performance with and without manipulation aids

Scene Index	Occupancy Density (bricks per m^2)	Approach	Failed Grasps	Double Grasps		Failed Pickups	Total Number of Failures	Successfully Sorted (out of 24)
				Same color	Different color			
1	243.9	Naive Method	7	0	3	3	13	21
		Our Method	1	0	0	0	1	24
2	404	Naive Method	8	2	2	3	15	21
		Our Method	1	1	0	2	4	22
3	574.2	Naive Method	4	2	1	6	13	18
		Our Method	1	1	0	0	2	24

TABLE I: Different kinds of failures for different scenes using the two approaches. The naive method does not take the state of the region into account. Sorting using manipulation primitives first manipulates a region based on its degree of clutter and then sorts it.

around p along the X axis because it would be easier for the pile to tumble in that direction (see Fig. 5b for an example). This action is likely to decompose the pile such that all Duplos lie directly on the tabletop.

After a region has been acted upon, the perception pipeline is invoked again and the whole procedure is repeated. Please see [10] for more details.

D. Results

We compared our approach to a ‘naive’ sorting approach where the perception pipeline essentially remains the same as described in section III-B but irrespective of the degree of clutter, the robot tries to pick and sort every brick without any manipulation primitives. Our evaluation metric is the number of successful grasps. We categorize failures as the following:

- Failed grasp: A grasp is attempted and the gripper moves to a brick to grasp it but fails.
- Double grasp: The gripper grasps two bricks instead of just the one it intended.
- Failed pickup: A grasp is not attempted because a brick is out of reach (too far on the table or dropped on the floor). Such a brick is never picked up.

We define *occupancy density*, a measure of the *degree of clutter* in the scene. Let \mathcal{A} be the minimum bounding box of all bricks in the scene and let n be the number of bricks. Then, the *occupancy density* is given by n/\mathcal{A} . This gives us an estimate of how densely packed the Duplos are.

For uncluttered scenes, both approaches reduce to the same algorithm and are successful in sorting all Duplos correctly. Fig. 6a shows a relatively uncluttered scene of dimensions $0.41\text{ m} \times 0.24\text{ m}$, Fig. 6b shows a cluttered scene of dimensions $0.27\text{ m} \times 0.22\text{ m}$, and Fig. 6c shows a piled up scene of dimensions $0.22\text{ m} \times 0.19\text{ m}$, each containing 24 bricks.

Table I summarizes the results for all three scenes. We observe that the total number of failures is much lower with our approach. While our approach results in all Duplos being successfully sorted for the relatively uncluttered and piled up scenes (scenes 1 and 3), there are 2 failed pickups in the cluttered scene (scene 2). This happened because one of the spread actions resulted in pushing 2 Duplos out of reach of the right arm of the PR2. Since our current setup does not involve moving the robot base or using the robot’s left arm, these 2 Duplos could not be picked up and sorted. A video showing clips from our experiments is available at <http://youtu.be/s-NMZMQmWjw>.

E. Discussion

The Euclidean color-based clustering will occasionally fail to distinguish between two objects in contact when they are of the same color. This might result in multiple objects of the same color being detected as a single object and grasped together. In Table I, note that our method never resulted in a double grasp of different color but there were instances of double grasps of same color. To overcome this shortcoming, there are a few other features that can be incorporated in object clustering. Assuming that the kind of objects being sorted are known, feature-based template-matching could be employed for segmentation. A knowledge of the object sizes could also help in detecting multiple objects in contact. Assuming that household objects are mostly convex, any concave corners would also suggest a contact between two objects. On the grasp planning side, we could give higher preference to grasps that grasp an object from a shorter edge. All points forming a short edge are more likely to belong to the same object and such a grasp is likely to pick up a single object. Singulation technique as described in [3] could be helpful too.

We also observe that the number of successfully sorted Duplos (last column of Table I) is not very different for our method and the naive method. So where does our approach win? To understand this, note that the perceptual pre-processing takes ~ 6 seconds, a pick and place action ~ 24 seconds, a spread action ~ 7 seconds, and a tumble action ~ 10 seconds. Thus, manipulation is the slower component of the pipeline and manipulation-aided sorting incurs an additional time cost. However, for highly cluttered scenes, naive sorting results in many failed grasps and retries while our approach first manipulates the scene to make it less cluttered and then attempts any grasps. Thus, the number of successful grasps goes up resulting in fewer grasp attempts and thus, the additional overhead of manipulation primitives is compensated for. This makes our approach more suited for cluttered scenes.

It is also worth mentioning that although we tested the manipulation-aided sorting algorithm on Duplos, the same principles will apply to other tabletop objects too as long as they can be reasonably segmented using Euclidean color-based clustering or any other segmentation algorithm. Good segmentation is needed to define the regions and their states correctly and consequently, to manipulate the scene appropriately.

IV. ENVIRONMENT EXPLORATION IN CLUTTER

Let us now investigate the use of manipulation for environment exploration which, in turn, could be applied to object search or counting objects in small, cluttered environments e.g., shelf of a refrigerator. The robot can rearrange objects but is not allowed to remove them. Since the environment is partially observable, an observe-plan-act strategy is called for.

A. Algorithm

We make the following assumptions for solving this problem:

- 1) The size of the world is (approximately) known.
- 2) The robot has perfect object segmentation capability.
- 3) If searching for an object, the robot has perfect object recognition capability.
- 4) The environment is a grid-world with resolution equal to the largest object size. Thus, each object fits within a grid cell. Additionally, each cell contains at most one object.
- 5) An object is visible to the camera only if there is no object in front of it in the grid.
- 6) Objects are not in contact with each other.

This simplified setup is only a starting point to begin understanding the nature of this problem. We propose an adaptive look-ahead exploration algorithm that guarantees complete exploration of the environment unless there are no more actions possible. To start with, it assigns a state of free, occupied, or unknown to every grid cell. The goal is to know the state of every cell and only the objects that are visible at a particular instant can be moved by the robot. Allowed actions are to move a visible object from its current cell to one of the four neighboring cells that is not occupied. It takes a starting horizon value, h_0 , as input and in every iteration,

Algorithm 1 Adaptive Horizon Exploration(N_r, N_c, H)

```

1:  $V = \text{Find-Visible-Objects}$ 
2:  $Occ = \text{Update-Occupied-cells}$ 
3:  $Free = \text{Update-free-cells}$ 
4:  $N(\text{unknown}) = N_r * N_c - (Occ + Free)$ 
5: while  $N(\text{unknown}) > 0$  do
6:    $H_{curr} = H$            { $H_{curr}$  - current horizon}
7:    $[IG, A] = \text{plan}(V, Occ, Free, H_{curr})$ 
   { $IG$  - information gain,  $A$  - action sequence}
8:   if  $IG = 0$  then
9:      $H_{curr} = H_{curr} + 1$ 
10:    Re-plan. Go to step 7.
11:  end if
12:  for all  $a_i \in A$  do
13:    Execute action  $a_i$ 
14:    if Action failed then
15:      Re-plan. Go to step 21.
16:    else if Revealed a previously unseen object then
17:      Re-plan. Go to step 20.
18:    end if
19:  end for
20:   $V = \text{Find-Visible-Objects}$ 
21:   $Occ = \text{Update-Occupied-cells}$ 
22:   $Free = \text{Update-free-cells}$ 
23: end while

```

chooses the sequence of actions of length h_0 that would reveal most of the unknown cells amongst all action sequences. We refer to the number of unknown cells whose state is revealed by an action sequence as the *information gain* of that action sequence. If all h_0 length plans result in no information gain, the horizon is incremented by 1 until a non-zero information gain is obtained. The algorithm then plans with the longer horizon, executes the plan, and falls back to h_0 for the next iteration. While executing the planned action sequence, if an unexpected discovery is made (an object unseen hitherto is seen), the action sequence is terminated and a new plan is generated taking the new information into account. This procedure is repeated until all cell states are known (free or occupied).

The details of the algorithm are given in Algorithm 1. The inputs to the algorithm are the number of rows in the discretized world grid (N_r), number of columns in the grid (N_c), and the starting horizon value (H). The algorithm has the property that it chooses longer look-ahead, which requires more planning time, only when it is absolutely necessary and hence, it is adaptive.

If planning and manipulation are costly, we could also keep the horizon fixed at h_0 and explore as much of the world as we can without incrementing the horizon. This may result in partially explored environments but would save on the number of actions taken and time used for planning.

B. Simulation Results

We carried out simulations of Algorithm 1 in MATLAB. Fig. 7 depicts the steps of the algorithm for a 3×4 grid with 7 objects and an initial horizon of 2. This environment was fully explored in 3.35 seconds using 4 actions. This time,

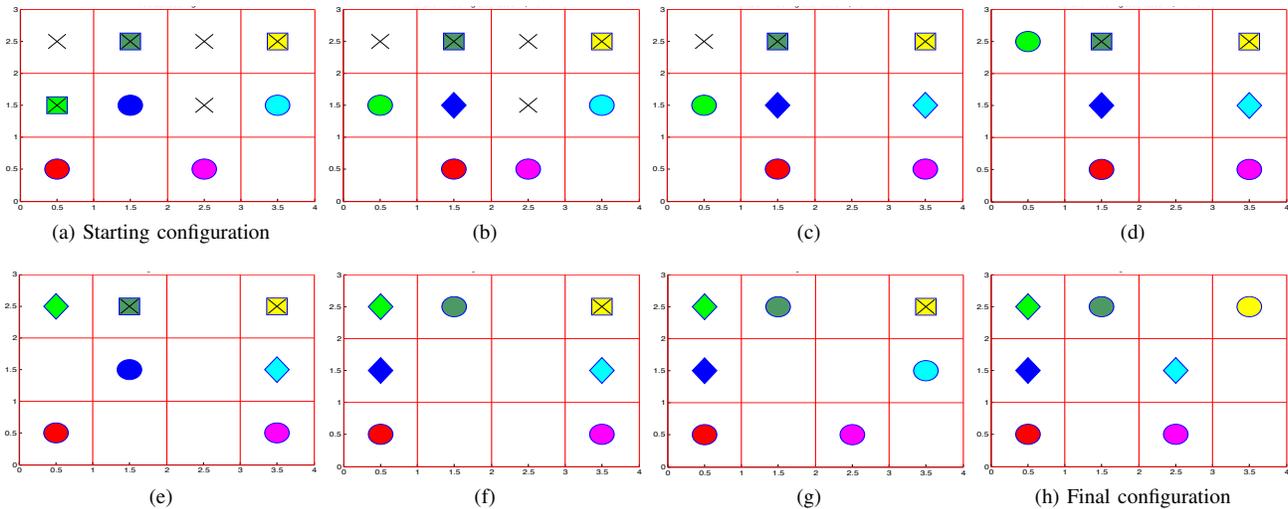


Fig. 7: Simulation of the adaptive horizon exploration on a 3×4 grid with 7 objects and a horizon of 2. The camera is placed along the lower boundary of the grid. Circles represent objects that are visible, squares represent objects that have never been seen so far and are thus, unknown, diamonds represent objects that have been seen at least once so far and so, their locations are exactly known, and crosses represent grid cells whose state is unknown. The cells with no object and no cross are known to be free. In the final configuration, all cell states are known, as shown by the absence of cross marks. This simulation took a total of 4 planning iterations, 7 actions, and 3.35 seconds.

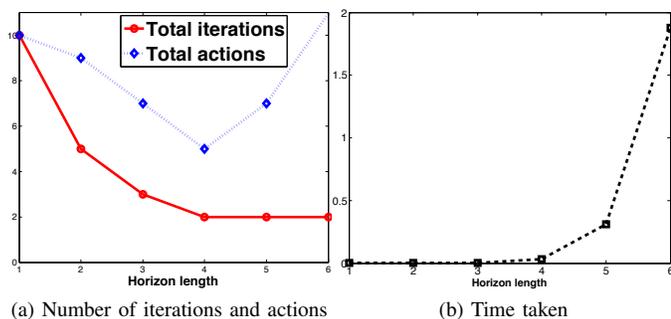


Fig. 8: Comparison of the performance of the adaptive horizon exploration with varying horizons.

of course, does not include any time required for the actual manipulation of objects.

Fig. 8 shows the effect of starting with different horizon values on the performance. The test environment was a 3×5 grid with 6 objects. We observe that though the number of planning iterations goes down as horizon increases, each iteration takes longer and in fact, the number of actions required does not decrease monotonically. Similar tests on many other environments of similar size suggest that a horizon value of 2 or 3 is often the best choice.

We also compare our approach to a random planning approach where a sequence of actions is chosen randomly from the available set of actions at any time. Results comparing the performance of the two approaches for a 3×5 grid-world and a horizon of 2 with increasing clutter are shown in Fig. 9. We plot the ratio of the average (over 50 different placements of objects) number of actions taken for random planning to those for adaptive look-ahead planning. We see that random planning needs 2 to 6 times more actions than our approach before the whole grid is explored. Since manipulation is

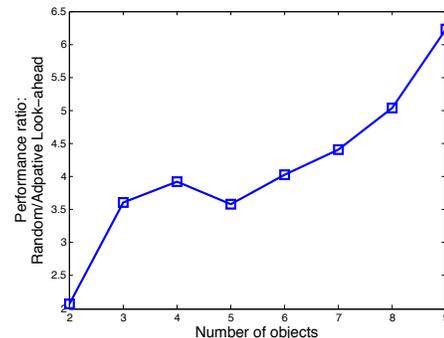


Fig. 9: Comparison of the performance of the adaptive horizon exploration with random planning. X-axis shows the number of objects in a 3×5 grid-world and the Y-axis shows the performance ratio for the random planning algorithm to the adaptive look-ahead exploration.

the slower component on an actual robot, exploration using random planning would clearly be much slower, particularly as the degree of clutter increases.

C. Discussion

We are now implementing this algorithm on the PR2 to test its feasibility in the real world. In practice, the assumptions of deterministic actions and perfect object segmentation are unrealistic. Some actions will fail or will be executed only partially. Similarly, object segmentation may become erroneous when objects are in contact with each other. The next natural step, therefore, is to model this uncertainty using a probabilistic planner. It will also be interesting to choose from a library of action primitives (as in Duplo sorting) depending on the kind of object being displaced and its surrounding environment. For example, it may be possible to pick and place an object that is not surrounded closely by other objects. However, in more clutter, it may only be possible to push and slide objects

around. Algorithm 1 takes actions based only on the cell states and thus, its planning time is independent of the robot. However, if there is a library of action primitives to choose from, the DOF of the manipulator arm will also affect planning time.

V. CONCLUSION

This paper explored the idea of manipulation as an aid to perception, and consequently, to the overall robotic manipulation task in cluttered environments. We studied its utility in two contexts - sorting small objects on a tabletop and exploring small cluttered environments. We presented a robust pipeline that uses two simple motion primitives to manipulate a cluttered scene in ways that help the robot to improve its perception - directly resulting in more successful grasps and more accurate sorting of Duplo bricks by color or size. Experimental results on the PR2 comparing sorting with and without the aid of manipulation suggest the benefits of the former, particularly as the degree of clutter increases. We also presented preliminary simulation results for an environment exploration algorithm that interleaves adaptive look-ahead planning with object manipulation. This has direct applications in object search, object counting, and scene mapping.

REFERENCES

- [1] A. Aydemir, M. Göbelbecker, A. Pronobis, K. Sjö, and P. Jensfelt. Plan-based object search and exploration using semantic spatial knowledge in the real world. In *Proc. of the European Conference on Mobile Robotics*, 2011.
- [2] K.F. Bohringer, K. Goldberg, M. Cohn, R. Howe, and A. Pisano. Parallel microassembly with electrostatic force fields. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, volume 2, pages 1204–1211, 1998.
- [3] L. Chang, J.R. Smith, and D. Fox. Interactive singulation of objects from a pile. In *Proc. of the IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, 2011.
- [4] B.J. Cohen, S. Chitta, and M. Likhachev. Search-based planning for manipulation with motion primitives. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, pages 2902–2908, 2010.
- [5] M.R. Dogar and S.S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. In *Proc. of the IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, pages 2123–2130, 2010.
- [6] P. Fitzpatrick. First contact: an active vision approach to segmentation. In *Proc. of the IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, volume 3, pages 2161–2166, 2003.
- [7] P.E. Forssén, D. Meger, K. Lai, S. Helmer, J.J. Little, and D.G. Lowe. Informed visual search: Combining attention and object recognition. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, pages 935–942, 2008.
- [8] B.P. Gerkey, S. Thrun, and G. Gordon. Visibility-based pursuit-evasion with limited field of view. *The Intl. Journal of Robotics Research*, 25(4):299–315, 2006.
- [9] S. Gould, P. Baumstarck, M. Quigley, A.Y. Ng, and D. Koller. Integrating visual and range data for robotic object detection. In *ECCV workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications (M2SFA2)*, 2008.
- [10] M. Gupta and G. S. Sukhatme. Using Manipulation Primitives for Brick Sorting in Clutter. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, 2012.
- [11] H.Y. Jang, H. Moradi, P. Le Minh, S. Lee, and J.H. Han. Visibility-based spatial reasoning for object manipulation in cluttered environments. *Computer-Aided Design*, 40(4):422–438, 2008.
- [12] D. Kang and K. Goldberg. Sorting parts by random grasping. *IEEE Transactions on Robotics and Automation*, 11(1):146–152, 1995.
- [13] D. Katz and O. Brock. Manipulating articulated objects with interactive perception. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, pages 272–277, 2008.
- [14] S. Koenig and M. Likhachev. D* lite. In *Proc. of the National Conference on Artificial Intelligence*, pages 476–483, 2002.
- [15] D. Kragic, M. Björkman, H.I. Christensen, and J.O. Eklundh. Vision for robotic object manipulation in domestic settings. *Robotics and Autonomous Systems*, 52(1):85–100, 2005.
- [16] V. Kumar, D. Rus, and S. Singh. Robot and sensor networks for first responders. *Pervasive Computing*, 3(4):24–33, 2004.
- [17] W.H. Li and L. Kleeman. Segmentation and modeling of visually symmetric objects by robot actions. *The Intl. Journal of Robotics Research*, 30(9):1124–1142, 2011.
- [18] M.L. Littman, A.R. Cassandra, and L. Pack Kaelbling. Learning policies for partially observable environments: Scaling up. In *Machine Learning-Intl. Workshop then Conference-*, pages 362–370, 1995.
- [19] J. Ma, T.H. Chung, and J. Burdick. A probabilistic framework for object search with 6-dof pose estimation. *The Intl. Journal of Robotics Research*, 30(10):1209–1228, 2011.
- [20] T. Parsons. Pursuit-evasion in a graph. *Theory and applications of graphs*, pages 426–441, 1978.
- [21] Radu Bogdan Rusu. *Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments*. PhD thesis, Computer Science department, Technische Universität München, Germany, October 2009.
- [22] R.B. Rusu, S. Cousins, and W. Garage. 3D is here: Point Cloud Library (PCL). In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, 2011.
- [23] D. Schiebener, A. Ude, J. Morimoto, T. Asfour, and R. Dillmann. Segmentation and learning of unknown objects through physical interaction. In *11th IEEE-RAS Intl. Conference on Humanoid Robots (Humanoids)*, pages 500–506, 2011.
- [24] C. Thorpe, J. Carlson, D. Duggins, J. Gowdy, R. MacLachlan, C. Mertz, A. Suppe, and B. Wang. Safe robot driving in cluttered environments. *Robotics Research*, pages 271–280, 2005.
- [25] C. Tovey, S. Greenberg, and S. Koenig. Improved analysis of D*. In *Proc. of the IEEE Intl. Conference on Robotics and Automation*, volume 3, pages 3371–3378, 2003.
- [26] Y. Ye and J.K. Tsotsos. Sensor planning for 3d object search. *Computer Vision and Image Understanding*, 73(2):145–168, 1999.