

Using Manipulation Primitives for Brick Sorting in Clutter

Megha Gupta and Gaurav S. Sukhatme

Abstract—This paper explores the idea of manipulation-aided perception and grasping in the context of sorting small objects on a tabletop. We present a robust pipeline that combines perception and manipulation to accurately sort Duplo bricks by color and size. The pipeline uses two simple motion primitives to manipulate the scene in ways that help the robot to improve its perception. This results in the ability to sort cluttered piles of Duplo bricks accurately. We present experimental results on the PR2 robot comparing brick sorting without the aid of manipulation to sorting with manipulation primitives that show the benefits of the latter, particularly as the degree of clutter in the environment increases.

I. INTRODUCTION

Imagine searching for a pen on your untidy desk. You look at the desk but can't see a pen. What is the next thing you do? You pick up and move other objects on the desk until you find the pen. The objects that you were not interested in got displaced in the process but that is acceptable. Using manipulation to better perceive the environment comes very naturally to us. We use it all the time not only for searching but also for mapping our environment. Shuffling through the pieces of a jigsaw puzzle, moving aside clutter to see what lies beneath - the examples are endless. An appealing trait of the process (other than its effectiveness) is that the manipulation maneuvers used are few and simple eg. picking up an object, pushing things aside, and shuffling through them.

This motivates our research in using manipulation of the environment as an aid to perception, and thus ultimately to purposeful grasping and manipulation. Humanoid and semi-humanoid robots will be a large part of future households to assist humans in various daily chores like cleaning, cooking, and taking care of the elderly. This necessitates a capability to understand their environment in real-time with high accuracy. Improving their perceptual abilities is only one side of the coin. Their ability to manipulate the environment can go a long way in overcoming the shortcomings of perception.

In this paper, Duplo bricks scattered on a table have to be sorted by a robot into different bins by size and color using data from a head-mounted Kinect sensor. We demonstrate through experiments that better 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. This results from the fact that an object that is graspable under uncluttered conditions (a Duplo brick in this case) becomes ungraspable as clutter around it increases. Clutter may also obscure some objects

and lead to incorrect perceptual results. We use simple manipulation primitives to reduce clutter and thus, aid the sorting process.

II. RELATED WORK

Several efforts have been directed at successful manipulation in cluttered spaces. Cohen, Chitta, and Likhachev [1] use heuristics and simple motion primitives to tackle the high-dimensional planning problem. Jang et al [2] present visibility-based techniques for real-time planning for object manipulation in cluttered environments. Hirano and Kitahama, and Yoshizawa [3] describe an approach using RRTs to come up with grasps even in cluttered scenes. Dogar and Srinivasa introduced the concept of *push-grasping* which manipulates and pushes obstacles aside to plan a more direct grasp for an object [4].

At the other end of the spectrum, the problem of improving perception in cluttered environments has also been studied. Thorpe et al [5] present a sensor-fusion system to monitor an urban environment and safely drive through it. Various techniques using multiple views, sensors, and computer vision algorithms have been developed to improve perceptual abilities of a robot in clutter ([6], [7], and [8]).

However, these two classes of work focus on either accurate grasping or object recognition and consider collisions with obstacles unacceptable. In a similar vein, there has also been research on achieving high perception accuracy so that grasp planning can benefit from it. These approaches study perception techniques to aid manipulation while we are aiming at the converse.

There has been some work on manipulation-aided perception, also referred to as *interactive perception* in literature. Using manipulation for object segmentation and feature detection has been studied by several groups ([9], [10], [11], [12]). Katz and Brock [13] use manipulation to build a kinematic model of an unknown object. Chang et al [14] show how objects could be separated from clutter using physical interaction.

In this paper, we have used Duplo sorting as an example to show that manipulation can improve robotic perception and grasping in clutter. Parts sorting is an old field of research and highly accurate vibration-based parts feeders have been designed in industry ([15], [16]). A parts feeder can not only sort but also position and orient bulk parts before they are fed to an assembly station and thus, help in micro-assembly. Work has also been done on sorting parts by shapes using Bayesian techniques and parallel-jaw grippers ([17], [18]). However, these methods assume availability of very specific

The authors are with the Robotic Embedded Systems Laboratory and the Computer Science Department, University of Southern California, Los Angeles, CA, USA megthagup@usc.edu, gaurav@usc.edu

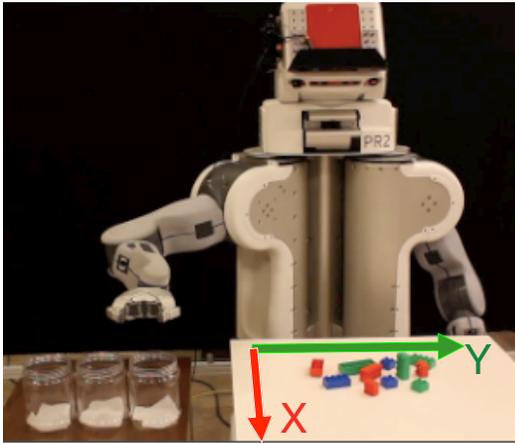


Fig. 1: Experimental setup. The PR2 with a head-mounted Kinect views Duplo bricks on a flat table-top. The task is to sort the bricks by color and size into the three bins placed next to the table.

equipment or require parts to be fed to the feeder one by one. Our method does not have these requirements.

III. EXPERIMENTAL SETUP

We carried out all our experiments on the PR2 robot, a semi-humanoid robotic platform developed by Willow Garage. PR2 uses the ROS operating system [19], an open-source operating system for robotics. The PR2 is equipped with a variety of sensors which include a stereo camera pair in its head, a tilting Hokuyo laser scanner above its shoulders, and cameras in its arms. We mounted a Microsoft Kinect sensor on the head. All perception data in this paper are from this sensor. The PR2 has a simple two-finger gripper using which it can execute only simple parallel-fingered grasps. This can interfere with achieving a robust grasp when there is not enough space on both sides of the object to be grasped. Additional manipulation is, thus, often important before grasping to ensure a successful grasp.

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 and Y axes as shown in Fig. 1.

IV. ALGORITHM

For the setup described in Section III, 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 V). 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*: Each Duplo brick in the region is not in contact with any other bricks.
- *Cluttered*: There are bricks that are close to or in contact with at least one other brick. However, the bricks are not piled up. All of them lie directly on the table.
- *Piled*: The bricks are cluttered and piled up, i.e., some of them lie on top of other bricks.

Fig. 2 depicts what each state might look like.

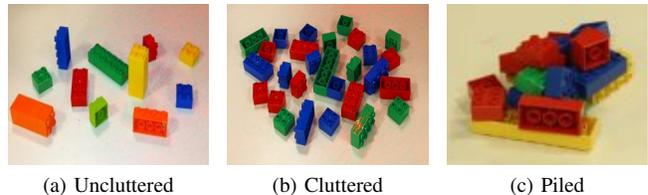


Fig. 2: The three different states a set of duplos could be in.

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.

The robot uses a minimalist library of manipulation primitives to reach a perceptual state where it can easily distinguish between different bricks and knows how to grasp them successfully. The end goal is accurate sorting with no failed grasps. The implementation details of the algorithm follow in the next section.

V. IMPLEMENTATION

A. Perception

The perception pipeline receives a point cloud from the head-mounted Kinect sensor and processes it using the following steps:

- 1) *Object Cloud Extraction*: Standard algorithms available in the Point Cloud Library (PCL) [20] are used for preprocessing. The point cloud is first sent to a pass-through filter that also takes in XYZ limits as arguments and retains only those points that lie within those limits. Assuming all objects are on a table and the table is placed in front of the robot, approximate XYZ limits are given and only the point cloud for the table and the tabletop objects are filtered out. Since exact location and dimensions of the table are assumed to be unknown, some outliers may still be present. Planar segmentation is employed to extract the largest plane (i.e. the table) from the point cloud. Dimensions of this plane give us the extent of the table and pass-through filtering is used once again with the table extents defining the limits this time around. This produces the correct point cloud associated with all the tabletop objects.

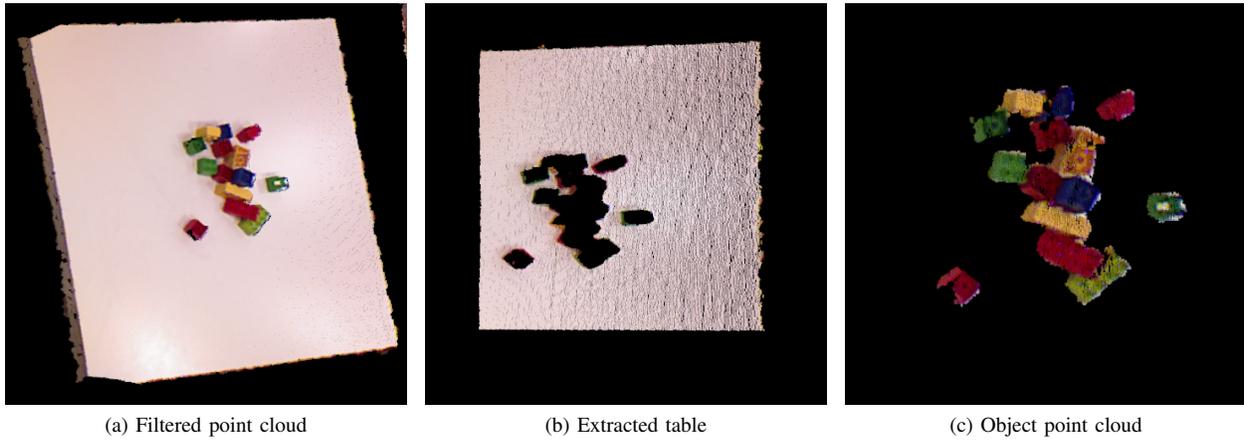


Fig. 3: Object cloud extraction from the point cloud received from Kinect.

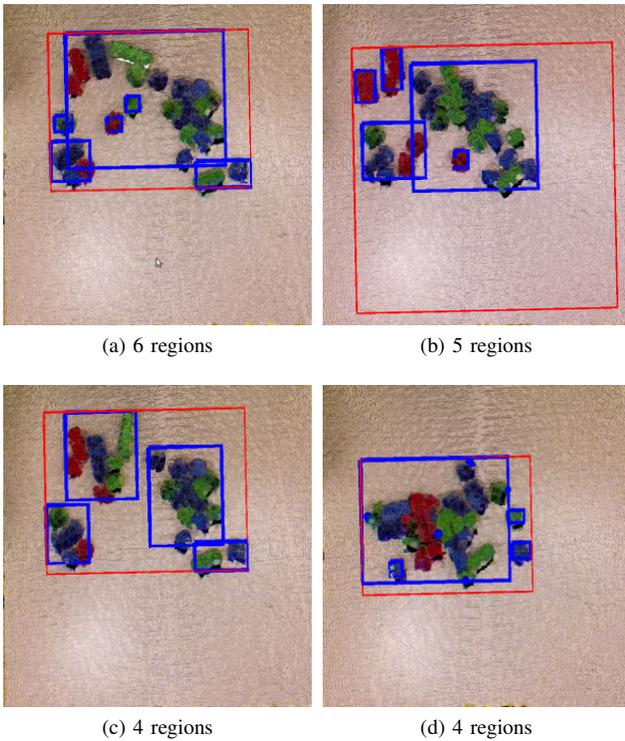
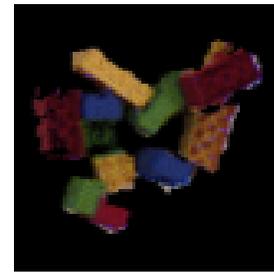


Fig. 4: Snapshots from RViz showing the different regions generated by the perception pipeline. Bricks within a region are very close to each other while any two bricks in different regions are a minimum distance apart. The red rectangle shows the bounding box for the robot’s estimate of the brick locations. The blue rectangles show the minimum bounding boxes for each region.

2) *Dividing into Regions*: The object point cloud is processed to obtain different regions such that the objects within a region are all closer than a pre-defined threshold to each other while objects in different regions are at least as far as the threshold. The Euclidean clustering algorithm available in PCL [21], which clusters the point cloud based on Euclidean distances, is used for this. Fig. 4 shows the regions generated by the algorithm for different scenes.

Learning the number of regions from the point cloud



(a) Object point cloud



Fig. 5: Clusters obtained from Euclidean color-based clustering.

makes the algorithm adaptive and allows the best manipulation decisions to be taken irrespective of the distribution of the bricks on the table.

3) *Assigning States to Regions*: For each region, the point cloud is processed to obtain meaningful clusters corresponding to individual Duplo bricks. To achieve this, the Euclidean clustering algorithm is modified slightly to take into account the color of the clusters as well. Since every Duplo brick is a single-color object, we look for neighboring points of the same color. If such a set of points satisfies the criteria of the minimum and the maximum number of points in a cluster, we obtain a cluster possibly corresponding to a single Duplo brick. Euclidean color-based clustering allows two bricks of different colors to be distinguished from each other even if they are in contact and their point clouds overlap.

Fig. 5 shows the result of Euclidean color-based cluster extraction for the object cloud shown. We see that Euclidean color-based clustering gives good results even in the presence of clutter. Note that our assumption of single-colored objects

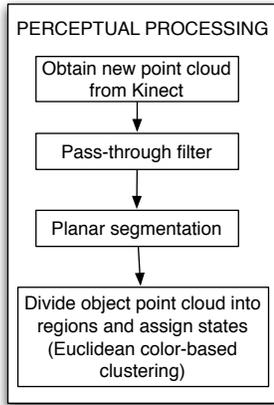


Fig. 6: The perception pipeline

is a way to simplify object segmentation since that is not our focus. For more complex objects (including multi-colored objects), we could employ a standard object segmentation algorithm and our technique of reducing clutter using manipulation primitives would still work.

For each brick in a region, we compute the number of bricks it is very close to or in contact with. Over all bricks in a region, the maximum value of this number is computed. Based on this number and the height of the point cloud, the region is classified as *uncluttered*, *cluttered*, or *piled*.

Fig. 6 shows the complete perception pipeline.

B. Manipulation Primitives

1) *Pick and Place*: For an uncluttered region, a slightly modified PR2 tabletop manipulation pipeline, which combines tabletop detection, collision map building, arm navigation, and grasping and is available as a ROS package, is used to pick and place Duplos into bins. Since collision with other Duplo bricks while grasping a Duplo brick is acceptable, the collision map includes only the table but none of the objects on it. Also, collisions with the table while grasping and lifting up an object are allowed. The point cloud of the uncluttered region is sent to the *grasp server* which attempts to pick up every cluster in the 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, the point cloud is sent to the *spread server* which looks for the brick that is in contact with the maximum number of other bricks. Perturbing this brick from its original position moves the neighboring bricks as well. Thus such motions 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.

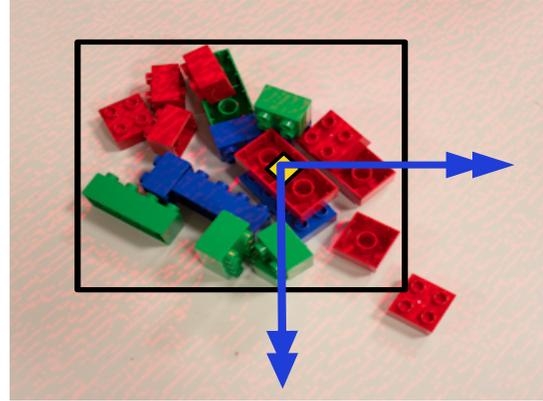
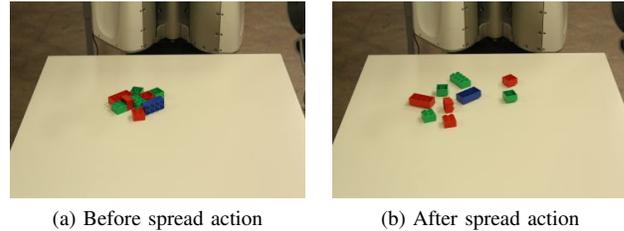


Fig. 7: Calculating directions to spread the region in.

For example, in Fig. 7, the black rectangle around the Duplos shows the minimum bounding box for that region and the red brick in the center is the one in contact with the maximum number of other bricks. The spread directions are chosen as the ones shown by blue arrows because the boundaries of the bounding box are closer in those directions. This action is likely to spread out the bricks making the region less cluttered than before.

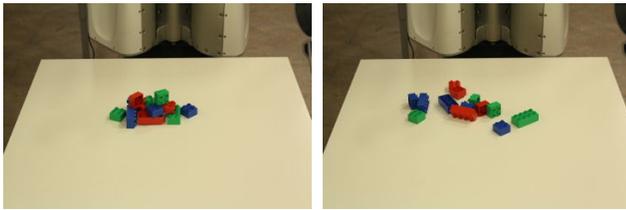


(a) Before spread action (b) After spread action

Fig. 8: Spreading results

3) *Tumble*: For a region containing a pile or multiple piles, the point cloud is sent to the *tumble server* which locates the point, p , at maximum height in the point cloud. Let us call 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 the movement is decided by the same method as in the spread action primitive (Fig. 7) and the hand moves in the direction of smaller extent of the point cloud around p because it would be easier for the pile to tumble on that side. This tumble action is likely to result in the region ending up in the cluttered state. If it fails because an inverse kinematic solution is not found, the robot tries to tumble the pile by moving its hand along the Y axis. If that fails too, the spread action is called for this region.

Fig. 8 and Fig. 9 show the result of the spread and tumble actions respectively after they have been invoked twice. Note that in our actual implementation, any bricks thus isolated in the previous iteration are first sorted and only then, the next set of spread and tumble primitives applied. This prevents



(a) Before tumble & spread actions (b) After tumble & spread actions

Fig. 9: Tumbling results

any neighboring isolated bricks from coming close (due to a spread or tumble action) and forming a new cluttered region again.

After a region has been acted upon, the perception pipeline is invoked again and the whole procedure is repeated. Fig. 10 gives an overview of the complete pipeline.

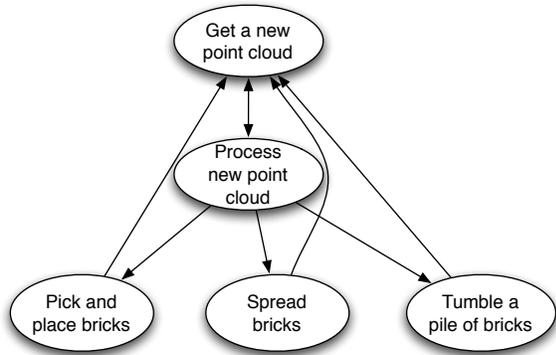


Fig. 10: A high level depiction of the Duplo sorting pipeline combining perception and manipulation.

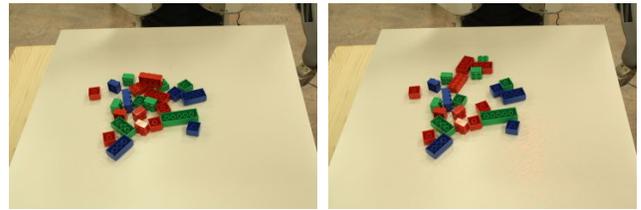
VI. RESULTS

Fig. 11 shows a sequence of snapshots of the table while the pipeline was running.

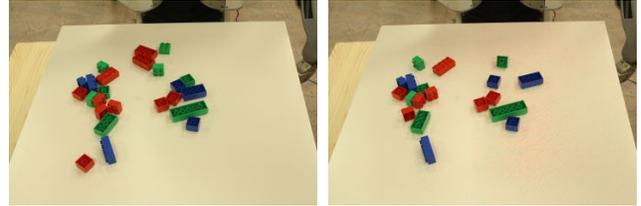
We compared our approach to a ‘naive’ sorting approach where the perception pipeline essentially remains the same as in Fig. 6 but the pick and place action is called for all regions irrespective of its state. Thus, the robot tries to pick and place every brick it sees, irrespective of the degree of clutter.

Our evaluation metric is the number of successful grasps. We categorize failures as the following:

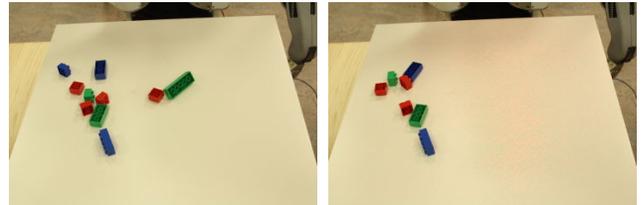
- Failed grasp: A grasp is attempted and the gripper moves to a brick to grasp it but fails.
- Double grasp: The gripper grasps two bricks instead of just the one it intended. This usually occurs when a brick is surrounded by other bricks and a non-optimal grasp is planned. Double grasp of bricks of same color may also occur because of the inability of the Euclidean color-based clustering algorithm to distinguish between them. Although a double grasp is inconsequential when sorting by color, it indicates a failure when sorting by size. It is also unacceptable for applications that require single object extraction.



(a) (b)



(c) (d)



(e) (f)

Fig. 11: Snapshots from the sorting experiment of a cluttered scene

- 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.

Clutter is a subjective term and what looks like a cluttered space to one person may be uncluttered to another. Therefore, 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} . Though this is only one of the many possible measures that could be defined for the degree of clutter, this gives us an estimate of how densely packed the Duplos are.



(a) Sorted by color (b) Sorted by length

Fig. 12: Sorting results

For uncluttered scenes, both approaches reduce to the same algorithm and are successful in sorting all Duplos correctly. Fig. 12 shows the results of sorting the bricks by color and length respectively for uncluttered scenes. For cluttered scenes, we present three examples comparing the performances of the two approaches, successively increasing the degree of clutter in the scenes.

| 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: Number of different kinds of failures for different scenes using the two approaches. The naive method does not take the state of the region into account. Our method first manipulates a region based on its degree of clutter and then sorts it.

- 1) Fig. 13 shows a relatively uncluttered scene of dimensions $0.41\text{ m} \times 0.24\text{ m}$ containing 24 bricks, giving an occupancy density of 243.9 bricks per unit area. The naive approach resulted in 13 unsuccessful grasps while our approach resulted in 1.

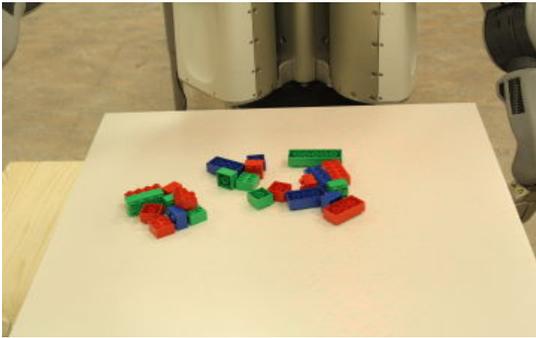


Fig. 13: Scene 1 - relatively uncluttered

- 2) Fig. 14 shows a cluttered scene of dimensions $0.27\text{ m} \times 0.22\text{ m}$ containing 24 bricks, giving an occupancy density of 404 bricks per unit area. The naive approach resulted in 15 unsuccessful grasps while our approach resulted in 4.



Fig. 14: Scene 2 - cluttered

- 3) Fig. 15 shows a piled up scene of dimensions $0.22\text{ m} \times 0.19\text{ m}$ containing 24 bricks, giving an occupancy density of 574.2 bricks per unit area. The naive approach resulted in 13 unsuccessful grasps while our approach resulted in 2.

Table I summarizes these results. We observe that our approach does better than the naive approach for all the scenes. In particular, the number of failures is much lower because our approach first manipulates the scene to make it



Fig. 15: Scene 3 - piled up

less cluttered (as shown in Fig. 11) and then attempts any grasps, thus increasing the likelihood of success.

Please see the attached video for a demonstration. The clips for the spread and tumble actions in the video are only snapshots of the complete sorting. At some places in the video, it may seem like bricks were removed by a person but they were actually picked up by the robot. Note that every time a set of spread and tumble primitives are applied in an iteration, a few bricks get isolated. These bricks are first picked up sorted, and only then a new set of spread and tumble actions are executed in the next iteration. However, these pick up actions are not shown in the video. Only when a brick was too far for the robot to reach or dropped on the floor, it was manually removed from the scene and counted as a failed pickup. It should also be noted that our manipulation primitives currently do not take the proximity of other regions into account while calculating spread and tumble directions. This could possibly result in an action bringing two regions closer and consequently, a longer sequence of manipulations. This can, however, be easily corrected and will be addressed in future work.

The naive sorting pipeline takes approximately 30 seconds to sort each brick - ~ 6 seconds for perceptual processing and ~ 24 seconds for grasp planning, moving the arm to the pickup location, picking the brick up, moving the arm to the bin location, and placing the brick. Thus, grasp planning and arm navigation are the components that slow down the pipeline. The manipulation-aided sorting adds manipulation of clutter and, hence, more instances of arm navigation to the pipeline, thus incurring an additional cost in terms of time. One spread action takes ~ 7 seconds while one tumble action takes ~ 10 seconds. However, for highly cluttered scenes, the naive sorting pipeline results in many failed

grasps and retries and consequently, more grasp planning, arm navigation and time overhead. The manipulation-aided sorting pipeline will compensate for the additional overhead of manipulation maneuvers by increasing the number of successful grasps in highly cluttered scenes.

VII. CONCLUSION

This paper explored manipulation-aided perception and grasping in the context of sorting small objects (Duplo bricks) on a tabletop. We presented a robust pipeline that combines perception and manipulation to accurately sort the bricks by color and size. The pipeline uses two simple motion primitives to manipulate the scene in ways that help the robot to improve its perception - directly resulting in more successful grasps. Preliminary experimental results on the PR2 robot and Kinect comparing sorting without manipulation to manipulation-aided sorting suggest the benefits of the latter, particularly as the degree of clutter in the environment increases. It is worthwhile to mention that our technique of reducing clutter using manipulation would be applicable even for more complex objects (including multi-colored objects) as long as there is a reasonably accurate object segmentation algorithm that we could use.

VIII. ACKNOWLEDGMENT

This work was supported in part by the US National Science Foundation Robust Intelligence Program (grant IIS-1017134), and by the Defense Advanced Projects Research Agency under the Autonomous Robot Manipulation program (contract W91CRBG-10-C-0136).

REFERENCES

- [1] B. Cohen, S. Chitta, and M. Likhachev, "Search-based planning for manipulation with motion primitives," in *IEEE Intl. Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 2902–2908.
- [2] H. Jang, H. Moradi, P. Le Minh, S. Lee, and J. Han, "Visibility-based spatial reasoning for object manipulation in cluttered environments," *Computer-Aided Design*, vol. 40, no. 4, pp. 422–438, 2008.
- [3] Y. Hirano, K. Kitahama, and S. Yoshizawa, "Image-based object recognition and dexterous hand/arm motion planning using rrt* for grasping in cluttered scene," in *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*. IEEE, 2005, pp. 2041–2046.
- [4] M. Dogar and S. Srinivasa, "Push-grasping with dexterous hands: Mechanics and a method," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2010, pp. 2123–2130.
- [5] 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*, pp. 271–280, 2005.
- [6] S. Gould, P. Baumstarck, M. Quigley, A. 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)*. Citeseer, 2008.
- [7] P. Forssén, D. Meger, K. Lai, S. Helmer, J. Little, and D. Lowe, "Informed visual search: Combining attention and object recognition," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*. IEEE, 2008, pp. 935–942.
- [8] D. Kragic, M. Björkman, H. Christensen, and J. Eklundh, "Vision for robotic object manipulation in domestic settings," *Robotics and Autonomous Systems*, vol. 52, no. 1, pp. 85–100, 2005.
- [9] P. Fitzpatrick, "First contact: an active vision approach to segmentation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 3. IEEE, 2003, pp. 2161–2166.
- [10] W. Li and L. Kleeman, "Segmentation and modeling of visually symmetric objects by robot actions," *The International Journal of Robotics Research*, vol. 30, no. 9, pp. 1124–1142, 2011.
- [11] E. Kuzmic and A. Ude, "Object segmentation and learning through feature grouping and manipulation," in *10th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2010, pp. 371–378.
- [12] D. Schiebener, A. Ude, J. Morimoto, T. Asfour, and R. Dillmann, "Segmentation and learning of unknown objects through physical interaction," in *11th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2011, pp. 500–506.
- [13] D. Katz and O. Brock, "Manipulating articulated objects with interactive perception," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2008, pp. 272–277.
- [14] L. Chang, J. Smith, and D. Fox, "Interactive singulation of objects from a pile," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2011.
- [15] K. Bohringer, K. Goldberg, M. Cohn, R. Howe, and A. Pisano, "Parallel microassembly with electrostatic force fields," in *IEEE International Conference on Robotics and Automation*, vol. 2. IEEE, 1998, pp. 1204–1211.
- [16] R. Taylor, M. Mason, and K. Goldberg, "Sensor-based manipulation planning as a game with nature," in *Proceedings of the 4th international symposium on Robotics Research*. MIT Press, 1988, pp. 421–429.
- [17] D. Kang and K. Goldberg, "Sorting parts by random grasping," *IEEE Transactions on Robotics and Automation*, vol. 11, no. 1, pp. 146–152, 1995.
- [18] A. Rao and K. Goldberg, "Shape from diameter," *The International journal of robotics research*, vol. 13, no. 1, p. 16, 1994.
- [19] "Robot Operating System (ROS)." [Online]. Available: <http://ros.org/wiki>
- [20] R. Rusu, S. Cousins, and W. Garage, "3D is here: Point Cloud Library (PCL)," in *Proceedings of the 2011 IEEE International Conference on Robotics and Automation, Shanghai, China*, 2011.
- [21] R. B. Rusu, "Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments," Ph.D. dissertation, Computer Science department, Technische Universität München, Germany, October 2009.