

Copyright © 2020 by Fan Wang
All rights reserved

AUTONOMOUS ROBOT PACKING OF COMPLEX-SHAPED OBJECTS

by

Fan Wang

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

Kris Hauser, Supervisor

Missy Cummings

Miroslav Pajic

Martin Brooke

Aaron Franklin

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Electrical and Computer Engineering
in the Graduate School of
Duke University

2020

ABSTRACT

AUTONOMOUS ROBOT PACKING OF
COMPLEX-SHAPED OBJECTS

by

Fan Wang

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

Kris Hauser, Supervisor

Missy Cummings

Miroslav Pajic

Martin Brooke

Aaron Franklin

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Electrical and Computer Engineering
in the Graduate School of
Duke University

2020

Abstract

With the unprecedented growth of the E-Commerce market, robotic warehouse automation has attracted much interest and capital investments. Compared to a conventional labor-intensive approach, an automated robot warehouse brings possible benefits such as increased uptime, higher total throughput, and lower accident rates. To date, developments in warehouse automation are mostly performed in aspects of inventory mobilization and object picking.

Recently, one emerging area that attracts a lot of research attention is the automated packaging or packing, a process during which robots stow objects into small confined spaces, such as shipping boxes. Automatic item packing is complimentary to item picking in warehouse settings. Packing items densely improves the storage capacity, decreases the delivery cost, and saves packing materials. However, it is a demanding manipulation task that has not been thoroughly explored by the research community. With the unprecedented growth of the E-Commerce market, robotic warehouse automation has attracted much interest and capital investments. Compared to a conventional labor-intensive approach, an automated robot warehouse brings possible benefits such as increased uptime, higher total throughput, and lower accident rates. To date, developments in warehouse automation are mostly performed in aspects of inventory mobilization and object picking.

Recently, one emerging area that attracts a lot of research attention is the automated packaging or packing, a process during which robots stow objects into small confined spaces, such as shipping boxes. Automatic item packing is complimentary to item picking in warehouse settings. Packing items densely improves the storage capacity, decreases the delivery cost, and saves packing materials. However, it is a demanding manipulation task that has not been thoroughly explored by the research

community.

This dissertation focuses on packing objects of arbitrary shapes and weights into a single shipping box with a robot manipulator. We seek to advance the state-of-the-art in robot packing with regards to optimizing container size for a set of objects, planning object placements for stability and feasibility, and increasing robustness of packing execution with a robot manipulator.

The three main innovations presented in this thesis are:

1. We are the first to implement a set of constraints into a packing planner, so that the resulting packing placements are stable and collision-free with a robot manipulator. Experimental evaluation of the method is conducted with a realistic physical simulator on a dataset of scanned real-world items, demonstrating stable and high-quality packing plans compared with other 3D packing methods.
2. Existing packing planner assumes the sequence of the packing can be controlled. However, in many cases, a controllable arrival sequence may not be practical. We formulate two variants of packing problems in which the set of items is known, but the arrival order is unknown. The goal is to certify that the items can be packed in a given container, and/or to optimize the size or cost of a container. That way, the items are guaranteed to be packable, regardless of arrival order.
3. We evaluate the proposed planner under real-world uncertainties such as vision, grasping, and modeling errors. We build hardware and software testbed that is fairly representative of current state-of-the-art sensing, perception, and planning for warehouse manipulation. A systematic evaluation of the testbed is then performed to study the sources of error and models their magnitude.

Robustness measures are proposed to improve the packing success rate under such errors.

Overall, empirical results demonstrate that up to 98% success rate can be achieved on a physical robot despite real-world uncertainties, demonstrating that these contributions have the potential to realize robust, dense automatic object packing in warehouses.

Acknowledgements

First and foremost, I would like to thank my advisor, Dr. Kris Hauser, for his continuous support and guidance throughout my Ph.D., and for providing me with the independence and freedom to work on a variety of research problems. I have been incredibly fortunate to work with Kris, who has been an excellent role model and inspired me to pursue a life-long career in robotics.

I would also like to thank my thesis committee members Dr. Missy Cummings, Dr. Martin Brooke, Dr. Pajic Miroslav, and Dr. Aaron Franklin, for their support and great advice. I am tremendously grateful for my friends and colleagues from the Duke Intelligent Motion Lab (Now the Intelligent Motion Lab at UIUC), including Yifan Zhu, Gao Tang, Shihao Wang, Joao Marcos Correia Marques and many others. It has been a great journal pursuing a Ph.D. with all of you, and I continue to learn andWith the unprecedented growth of the E-Commerce market, robotic warehouse automation has attracted much interest and capital investments. Compared to a conventional labor-intensive approach, an automated robot warehouse brings possible benefits such as increased uptime, higher total throughput, and lower accident rates. To date, developments in warehouse automation are mostly performed in aspects of inventory mobilization and object picking.

I want to thank the members of Team Duke when participating in the Amazon Robotics Challenge 2017 during my first year of Ph.D., Particularly Mark Draelos, Brenton Keller, Anne Draelos. Taking part in the challenge together was both a thrilling and humbling learning experience, and you all have inspired me to be a more exceptional Ph.D. student.

I am also very fortunate to have had the opportunity to get supported and work as an intern at Amazon Robotics with Jane Shi, an inspiring mentor, and friend.

During my internship, I have enjoyed working with many team members, including Joey Durham, Chaitanya Mitash, and Sachal Dillion.

Finally, I would like to thank my partner Haizhen Zhu for always being there and for putting my needs in front of his. I would also like to thank my parents and other family members for their unconditional love and support.

Contents

Abstract	iv
Acknowledgements	vii
List of Figures	xiii
List of Tables	xix
1 Introduction	1
2 Offline Bin Packing under Stable and Robot-feasible Constraints	9
2.1 Introduction	10
2.2 Problem Defination	11
2.2.1 Stability checking	12
2.2.2 Manipulation feasibility	13
2.3 Pipeline for Robot-packable Planning	13
2.3.1 Placement sequence	15
2.3.2 Generating ranked transforms	15
2.3.3 Pipeline summary and fall back procedures	17
2.4 Heightmap-Minimization Heuristic	18
2.5 Experiment	20
2.5.1 Robot manipulation feasibility with a vacuum gripper	21
2.5.2 Small Order Packing	22
2.5.3 Comparisons on Large Itemsets	23
2.5.4 Executing Packing Plans in Simulation	25
3 Robot Packing with Known Items and Nondeterministic Arrival Order	26

3.1	Introduction	27
3.2	Related work	29
3.3	Problem formulation	30
3.3.1	Constraint formulation	31
3.3.2	Nondeterministic problems	33
3.3.3	Container optimization variants	34
3.4	Method	34
3.4.1	Offline planner	35
3.4.2	Compatibility	35
3.4.3	Constraint dependency graphs	37
3.4.4	Coverage verification	38
3.4.5	Quasi-online packing	39
3.4.6	Analysis	42
3.5	Planning Heuristics	47
3.5.1	Dependency minimization heuristic	47
3.5.2	Matching prior placements	48
3.5.3	Container optimization heuristics	48
3.6	Experiments	49
4	In-hand Object Scanning via RGB-D Video Segmentation	53
4.1	Introduction	54
4.2	Related Work	56
4.3	In-hand Object Reconstruction Pipeline	58
4.3.1	BackFlow Video Segmentation Method	59
4.3.2	3D Reconstruction	64

4.4	Experiments	65
4.4.1	RGB-D in-hand object manipulation dataset	65
4.4.2	Novice scanning of many items	70
5	Systems and Analysis for Robust Robotic Packing	72
5.1	Introduction	73
5.2	Related work	74
5.3	System setup and error analysis	75
5.3.1	Experimental setup	75
5.3.2	Factors Affecting Packing Success	77
5.4	Error reduction methods	79
5.4.1	Open-loop packing baseline	80
5.4.2	Closed-loop packing	82
5.4.3	Robust planning	83
5.4.4	Closed-loop packing and robust planning	84
5.5	Analysis and experimentation	84
5.5.1	Itemset	85
5.5.2	Uncertainty evaluation	85
5.5.3	Summary statistic: placement error	86
5.5.4	Simulation testing	87
5.5.5	testing on physical platform	88
6	Robot Button Pressing In Human Environments	99
6.1	INTRODUCTION	100
6.2	RELATED WORK	102

6.3	CHARACTERIZATION OF SWITCHES IN HUMAN ENVIRONMENTS	103
6.3.1	Button Taxonomy	104
6.3.2	Operation	105
6.3.3	Location and Geometry	107
6.3.4	Force	107
6.3.5	Surface material	108
6.3.6	Travel distance	109
6.4	THE SWITCHIT PLATFORM	109
6.4.1	Hardware	109
6.4.2	Scallop fingertip	110
6.4.3	Environmental annotation and calibration	112
6.4.4	Button Panel Recognition and Localization	114
6.4.5	Button State Recognition	115
6.4.6	Control	116
6.5	EXPERIMENTS	117
6.5.1	Calibration	117
6.5.2	Measurement of system accuracy	118
6.5.3	Test panel experiments	119
6.5.4	Experiment pressing in an office building and home	121
6.5.5	Test panel experiments	123
7	Future Work	124
	Bibliography	127

List of Figures

1.1	Examples of poor space utilization in shipping boxes.	5
2.1	Example packing placements obtained by HM and DBLF. HM finds more compact and stable packing compared to DBLF.	20
2.2	(a) Grasp poses generated satisfying vacuum graspability constraints. (b) Compatible gripper poses with candidate object orientation are checked for clearance with the container and the object pile. Collision-free grasps are shown in green and colliding grasps are colored in transparent grey.	21
2.3	Distribution of the solution containers found under different level of constraints. The x axis is the 5 container dimensions tested.	22
2.4	Examples of packing plans for itemsets of size 10.	23
3.1	Feasible solutions for a 2D, 3-item instance of (a) NDOP and (b) QOP. All $3! = 6$ possible arrival orders are collision-free, loadable from top-down, and yield intermediate piles that are stable under gravity. In QOP, an item is never moved after it is placed.	28
3.2	Examples of plans that are infeasible for arrival order ABCD: (a) unstable, (b) items C and D collide with B along the loading direction, (c) and the path for the robot manipulator to grasp and load item D is infeasible. In (d), although ABCD is feasible, the prefix requirement is violated because the sub-plan ABC is unstable.	31
3.3	A plan and its dependency graph. C_2 requires D_1 to be present to maintain the stability constraint, because otherwise the imbalanced weight on B would cause tipping. Similarly, D_2 depends on C_1 , and so forth for C_3 and D_3 . This CDG is compatible with orders of the form $AB(C_1D_1)(C_2D_2)(C_3D_3)$ where the (XY) denotes either XY or YX	36
3.4	A set of plans P_1, \dots, P_4 (top) and their dependency graphs G_1, \dots, G_4 (bottom). For any ordering beginning with A, there is at least one plan (P_1 or P_2) compatible with it. But for any ordering beginning with BDA, CB, CD, DAC, or DC, no plans are compatible.	38

3.5	An example showing that QOP (Alg 8) is not necessarily complete even with a complete offline planner. (a) The first recursive call produces a feasible plan with A placed first. (b) Once item A is placed in the planned location, the plan is infeasible for order ACB, as shown in (c). On the other hand, if A was placed as in (d), a feasible QOP solution could result.	44
3.6	(a) Worst-case behavior of Verify-CDG-Coverage occurs in an instance with n plans, where $n - 1$ “books“ are stacked vertically with the n 'th book stacked horizontally on top. Every possible order of $k \leq n$ items is compatible with a set of $n - k + 1$ plans, and a recursion depth of n is required. (b) With a slightly different stacking, the dependency graphs are reversed. Only a depth 1 recursion is needed due to the singleton pruning step, so running time is polynomial.	45
3.7	Illustrating the reduction from SAT. Each clause (upper left) is converted into a dependency graph (lower left), and a counterexample (lower right) ordering corresponds to a SAT solution (upper right).	46
3.8	The dimensions (in inches) of containers 1–5 used in our experiments, in order of increasing length + girth.	48
3.9	An NDOP solution for packing 5 items in container 5.	50
3.10	Distribution of the solution container for offline / NDOP container optimization and various itemset sizes. Most small instances can be packed in any order, but with more items the variability in order requires larger boxes.	51
3.11	A failure case for packing 5 items in container 3. The offline planner solves for the arrival order in (a) but fails on the order in (b) because there is insufficient remaining space for the fifth item.	51
3.12	A QOP solution for 4 items in container 5. Due to the matching heuristic, only four plans are needed (one for each item placed last).	52
4.1	Some successfully reconstructed objects from a 200-item test set, scanned by a novice user of our system.	55

4.2	Skin regions identified by skin-pixel based segmentation described in [VSA03]. This method cannot detect shaded areas of hand as well as accessories on hand (jewelry, nail polish, etc.). It also falsely includes object regions with colors close to human hands.	58
4.3	The new sides problem poses a challenge for video segmentation algorithms that rely on coherent object appearance. Tracking a can of tomato soup with a hierarchical graph-based method [GKHE10] fails to grow the segmentation to newly appeared top of the can.	59
4.4	BackFlow takes user annotation of background or ground truth segment on the first frame and propagates labels as follows. 1) Initialize background labels (superpixels) from user input. 2) Propagate background labels from frame t to frame t+1 through optical flow. 3) Perform graph-based foreground derivation. 4) Remove background labels in segmented foreground and grow background labels to neighbourhood superpixels.	61
4.5	Background labels (marked as red dots) grow gradually into the previously unseen palm while maintaining separation from the foreground object.	64
4.6	RGB-D in-hand object manipulation dataset contains RGB-D video of 13 non-duplicate items from YCB food category, including Pringles, mustard, Cheez-It, sugar, Spam, tomato soup, banana, pear, plum, strawberry, orange, lemon and Jell-O.	66
4.7	Per sequence accuracy (IoU) comparison of BackFlow to other state-of-the-art video segmentation methods.	67
4.8	Reconstructed models from BackFlow segmentation results.	68
4.9	Comparison of reconstructions from OSVOS and BackFlow.	69
5.1	The experimental packing setup consists of a UR5 robot equipped with E-pick suction gripper, and overhead cameras (not pictured). Errors from pose recognition, camera-robot calibration, box location calibration, grasp planning, the center of mass estimation, and geometric modeling affect the overall success rate of packing.	91

5.2	Top-down diagram of the packing experiment. Objects to be packed are arranged in the picking area, and must be packed in the container (e.g., a shipping box). Overhead cameras (Realsense SR300 and Ensenso N35) provide color and depth information.	92
5.3	A diagram of how error sources (rectangles) contribute to unexpected events (rounded rectangles) and ultimately to packing failure. We use ϵ_{place} as a measure for grasp pose error.	93
5.4	Closing the loop around vision leads to more robust packing. (a) Offline plan. The 2 <i>cracker boxes</i> are stacked first, followed by <i>bleach</i> and 2 <i>soup cans</i> . (b) The second <i>cracker box</i> is shifted during packing, blocking the access for <i>bleach</i> . Open-loop packing would lead to a failure. (c) The closed-loop strategy captures a heightmap of the placed pile, detects a violation for the <i>bleach</i> placement, and plans a new location. (4) <i>Bleach</i> is placed successfully in the empty location.	94
5.5	Conservative geometry margins reduce inadvertent collisions. (a) A non-conservative plan (left) causes a failure during the execution of the plan (right). Specifically, the shifted <i>cracker box</i> caused a failure and crash of the <i>sugar</i> item planned on top of it. (b) A conservative plan with $\delta = 1$ cm margin. During execution, all items are placed within the container without causing an inadvertent collision.	95
5.6	15 items from YCB video object dataset are used in the packing experiment	96
5.7	Simulation packing success rates for 10-item orders. Pose error scaling factor is given on the horizontal axis, and the success rate is given on the vertical axis. The baseline (V1) success rate drops off dramatically as pose error increases, while the robust technique (V4) is much less sensitive to error.	97

5.8	<p>Example of failed packing executions, for each experimental condition. In the baseline (V1), all items are planned tightly, and the <i>cracker box</i> hit a container wall, causing it to tilt, and for all subsequent items planned on top of it to smash. In robust planning (V2), although margins were enforced in planning, a substantially incorrect pose estimation caused the <i>bleach</i> item to stick out of the container, and no action is taken to correct it. In closed-loop packing (V3), the first-placed <i>cracker box</i> caught on the side and failed to be pushed in. This was the only case when the push maneuver did not work in our examples. Although this caused a failed plan, replanning prevented the subsequent items from smashing onto the <i>cracker box</i>. With both closed-loop packing and robust planning (V4), a replan was triggered after the <i>cracker box</i> and before the <i>bleach</i> item. Because the pile was treated as a fixed, rigid object, the <i>bleach</i> was placed on top of the <i>cracker box</i>, but since it was unstable, this caused it to tip over. . . .</p>	98
6.1	<p>SwitchIt is a spherical robot equipped with an RGB-D camera. Button panels are annotated using a QR code sticker affixed during a manual setup phase. Shown here mounted on a tripod, the robot is preparing to press buttons on an electronic passcode panel.</p>	101
6.2	<p>Buttons characterized by type.</p>	106
6.3	<p>Cylindrical and scallop tip designs: (a) Scallop tip skeleton, (b) Scallop tip with coating, (c) Side view of tips. From top to bottom: scallop skeleton, cylindrical, scallop coated with rubber, (d) Top view of tips. From top to bottom: scallop skeleton, cylindrical, scallop coated with rubber</p>	111
6.4	<p>Steps of the panel calibration process: (a) Taking a picture of the panel panel with RGB-D camera, (b) Tap on a button to zoom in. Already calibrated buttons will be marked, (c) Draw rectangles in the area of interest as guided for each type of buttons, (d) Zoomed in button details make it easier for operators to calibrate with higher accuracy .</p>	113
6.5	<p>Point cloud and RGB view showing localization of button surface in real-time. Computed locations of the buttons are drawn as red rectangles.</p>	114
6.6	<p>Calibration areas by button type: (a) Push button, (b) Toggle, (c) Rocker, (d) Slider, (e) Turn knob, (f) Pull button.</p>	115


6.7	The above apparatus is used for testing the accuracy of the system: (a) Robot and camera are calibrated using a cross marker, (b) Center button is located at the intersection of opposing diagonal circles, (c) Center plate is removed during the test to not interfere with tip motion.	119
6.8	Histogram of system errors.	120
6.9	Error distribution in XY.	120
6.10	A selection of button panels that SwitchIt succeeded in activating: (a) Door passcode, (b) Disabled door exit, (c) Thermostat buttons, (d) Switch with dimmer, (e) Office stereo control, (f) Security door entrance, (g) Electronic keypad, (h) Light switch	122
6.11	Four button types that SwitchIt failed to activate: (a) An old-style timer, (2) An elevator pull button, (3) A washer control button, (4) A turn knob on oven.	123



List of Tables


2.1	Comparing planning techniques on 10-item orders with and without robot-packable constraints	24
2.2	Impact of Δr and candidate number N on the results	24
2.3	Execution success rates in simulation, 10-item orders	25
3.1	NDOP container optimization results	49
3.2	QOP planning results in container 5	51
4.1	Segmentation Results	67
4.2	Hausdorff Distance measurements w.r.t. bounding box size of the reconstructed models	70
5.1	Summary of empirically evaluated error sources	87
5.2	Success rates on the physical packing platform, 5-item orders	89
6.1	Breakdown of button characteristics by type	108
6.2	Test Panel Experiments	121

Chapter 1

Introduction

The two main factors that motivate warehouse automation are 1). to reduce the sheer size of a warehouse facility and to speed up transportation of goods within them, and 2). to automate labor-intensive and repetitive operations. 

Large areas are needed for warehouses to store the item  assortment in racks, to move stock, to unload and load trailers and containers, and to allow people to pick from them. Meanwhile, it becomes time-consuming to move items and people around inside such a large space. By earlier 2000s, a big boost of warehouse technology in address this problem is credited to AVS/R (autonomous vehicle-based or shuttle-based storage and retrieval) systems. The  systems use racks with aisles and deploy autonomous shuttles that operate to transport inventory around. Another important development to address this problem is the use of automated pallet stacking and destacking technologies.

Later in the 2000s, with advanced in sensor technologies and robotics, mobile robots such as automated guided vehicles(AGVs) and autonomous mobile robots (AMRs) are employed to move inventory shelves around in the warehouse and present inventory for the operator to retrieve . The inventory could include storage racks, packages, or a tote of items. Mobilized storage racks can be stored to a higher density as there is no need to leave room for aisle for a human to walk by. The fact that the robot brings the inventories to the operator instead of having the operator walk to them also increases the base pick rate. Compared to the earlier warehouse automation technologies that rely on mazes of conveyor belts and chutes, robotic automation allows for faster deployment and greater mobility.

The last ten years have seen a fast expansion of the E-commerce market. Many warehouses move to operate 24/7, handle large and variable daily order volumes, and store millions of unique items. The movements raise the demand for more comprehensive and sophisticated robotic automation. The hope is to fulfill a customer order fully automati-

cally, from sorting inbound inventory to pick and pack for each individual ordered item. Those processes are great candidates for automation: they are laborious and repetitive in nature, and often suffer from poor ergonomics. In fact, interest in robust and versatile robotic manipulation for a wide range of objects is almost as old as robotics. However, most automated robot success is restricted to known objects, in controlled configurations, with specialized hardware. The variety of the objects, varied widely in size and appearance, ranging from loose objects wrapped in plastic such as clothing, to soft objects like plush toys, as well as solid but highly concave objects like bowls and cups, present major challenges for robot vision and manipulation.

How robots can grasp and handle a wide range of objects robustly has long-term research focus of the robotic community. Of which the grasping is probably the most studied aspects. Grasping methods can be roughly categorized into model-based grasping and data-driven grasping. Model-based grasping selects a grasp primitive based on the geometry of the object, using criteria such as force closure [Ngu88] or caging [RMF12]. Data-driven methods sample grasp candidates for an object and rank the candidates according to a scoring metric. The metrics come from existing grasp experiences that can be a heuristic or an evaluation of the grasp generated in a simulation or on a real robot [BMAK14]. Progress over the past decade in machine learning, perception, and hardware design has enabled robots to grasp an expanding set of objects and to do so more robustly. While different research groups use the grasping dataset of various difficulty and evaluation metrics, one of the most state-of-the-art systems is arguably the Google grasping project [LPKQ16], which allows industrial robot arms to grasp novel objects in a cluttered container. A large convolutional neural network was trained, whose input is an image of the current scene, to predict the probability of grasping motions to result in successful grasps. Empirically it shows that the robot has an 80% success rate among 100 grasping attempts.

The Amazon Robotics Challenge (ARC) [Rut15] held yearly from 2015-2017 has drawn wide attention from the robotics community. The competition expands the research horizon from just grasping to purposely pick and place, and highlights the importance of the

complectory of a picking action: What comes after picking? How can objects be placed or stowed for different purposes of operation? Specifically, the rule of the competitions tests the ability of robotic systems to fulfill an order by autonomously picking the ordered items from a tote and place it in a storage unit, and pick the item from storage and pack it into a shipping box, which is reasonably good representation of the key challenges in a warehouse. Many state-of-the-art pick-and-place systems such as [YFD⁺16, SLMG⁺18, ZSY⁺18] have resulted from the competition. Those systems are capable of recognizing and grasping both known and novel objects to a degree in relatively cluttered environments and place them in over-sized storage space.

Following ARC, robot packing has attracted greater research interest. Robot packing studies the placing aspects of the manipulation, under a more rigorous problem definition of dense object packing. While APC has laid out some foundations for building a pick and place system, the placing aspects in such systems are largely simplified by using over-sized containers. Under this setting, packing can be addressed with simple heuristics or even dropping items from the top center. Simple heuristic methods are insufficient for addressing a dense object packing problem, during which robots stow objects into confined spaces to high density, such as shipping boxes, totes, or shelves.

To date, most of the dense packing has been studied as inserting a single object into a cluttered environment using high-resolution sensors. For example, Yu et al. [YR18] propose an insertion strategy that leverages tactile sensing to probe the gap for the designated insertion pose while monitoring incipient slip to maintain a stable grasp on the object, while estimating and correcting for residual position uncertainties. Shome et al. uses RGB-D data and a vacuum-based end-effector as a grasper and push finger. They monitoring of potential failures in real-time and uses corrective pushing to achieve tight packing of an object to its designed spot. Inserting an object into a tight space has its roots in the classic peg-in-hole problem, which a typical contact manipulation task that requires precise position and some form of compliance, either passive or active. Many approaches have been proposed to solve the problem. One class of methods is based on passive compliance hardware and

control algorithms. Drake [Dra05] designed a passive compliance device called Remote Center Compliance (RCC) for correcting small uncertainties in the assembly task. Park et al. [HJJ⁺13] proposed a strategy that adopted hybrid force/position control and passive compliance control for successful peg-in-hole assembly. Another class of methods is model-based active sensing, which utilizes feedback from sensors to identify the configurations or errors in the assembly process. The model-based methods are usually more adaptable to new environments. An example includes Bruyninckx et al. [BDD95], who proposed a model-based method to model different contact situations and deployed the model and feedback from a force sensor to explicitly find the hole and align the axes of the peg with the hole. Yu’s work also belongs to this class of method.

This dissertation studies dense packing, with a focus on a specific problem: pack-objects of arbitrary shapes and weights into a single shipping box. This problem poses additional challenges on top of the dense packing problem discussed (which is only concerned with inserting an object into its designated spot). The additional challenge is to optimize a placement plan and a container for a given set of objects, so that all objects in the set are guaranteed to fit within the container when the packing is performed by a robot.

The current state of practice in fulfillment centers leaves the responsibility of container selection and packing to human worker intuition. Due to demanding schedules, workers cannot employ much foresight in the packing process and are reluctant to re-pack. As a result, over-sized containers are often used, incurring waste and high shipping costs^{1.1}. Better containers and packing plans could be chosen using automated algorithms, whether packing is accomplished by humans or robots.

Traditionally, problems that involve the placement of objects within a container or a set of containers are referred to as cutting and packing problems. Most existing packing algorithms apply to idealized scenarios, such as rectilinear objects and floating objects not subject to the force of gravity. Popular variations of the cutting and packing problem include the bin and strip packing problem, the knapsack problem, the container loading problem, and others. Under specific settings, such problems can be formulated and solved

optimally using the exact algorithms. One example of these state-of-the-art exact algorithms is the solution to the 3D bin packing problem using branch and bound, proposed by Martello et al. [MV98, MPV00], whose work is further extended by many including Boef et al. [dBKM⁺05] and Crainic et al. [CPT08]. Heuristic methods and metaheuristic approaches have also been developed over the years, such as the popular Bottom-Left heuristic [BCR80] and the Best-Fit-Decreasing heuristic [JDU⁺74].

More recent work has addressed irregular shape packing, and only heuristic methods are practical because the search space is infinite. Metaheuristics are commonly used in this setting such as Simulated Annealing (SA) [kam88, ZH04, LLCY15] and Guided Local Search (GLS) [FPZ03, Ege09, VT03, VVHS15, BLP13] that start with an initial placement and iteratively improve the placement by moving the pieces in the neighborhood while minimizing an objective function (e.g., overlap in the system). Other work has also proposed constructive positioning heuristics for 3D irregular objects, such as Deepest-Bottom-Left-Fill (DBLF), which places items in the deepest, bottom-most, left-most position; and Maximum Touching Area (MTA), which places an item in a position that maximizes the total contact area of its faces with the faces of other items [WGC⁺10].

To perform automatic packing in warehouses using a pre-computed packing plan, several real-world issues need to be addressed, such as stability under force of gravity, and kinematics and clearance issues for the robot. If stability is not considered, the pile may shift during execution, and therefore subsequent placements are unlikely to be executed as



Figure 1.1: Examples of poor space utilization in shipping boxes.

planned. If kinematics and clearance are not considered, the robot may be asked to perform infeasible motions (e.g., grip an item from underneath, bring an item to a target through another interlocked item, or pass through the container wall). Unstable placements and infeasible robot motions may result in failure to contain all items, and even cause damages to the robot and items. A failed packing is usually difficult and time-consuming to recover from, e.g., remove all placed items in the old box and replace them in a larger box.

Some research has taken aspects of stability into consideration during packing. Egeblad et al., for example, use a two-stage GLS packing algorithm that, in the first stage, optimizes for the center of gravity and inertia of the pile and, in the second stage, minimizes overlap in the system [Ege09]; Liu et al. propose a constructive method that packs irregular 3D shapes using a Minimum-Total-Potential-Energy heuristic [LLCY15]. This method performs a grid search for the lowest gravitational center height Z for each placement. However, both proposed methods use heuristics only and do not verify the stability of each placement. In contrast from these works, our method enforces stability explicitly using constraints.

We also know of one packing work that takes into account robot manipulation feasibility [dBKM⁺05], in which the author proposes a variant of the orthogonal 3D box packing scheme such that no prior packed box is in front of, to the right of, or above the current placing box, to avoid possible collision with a vacuum gripper. This placing rule, however, cannot guarantee collision-free placements with other gripper geometries and neglects robot kinematic constraints and graspability constraints.

We are the first to propose a comprehensive set of constraints tailored to the automated warehousing domain and implement a packing algorithm to solve those constraints. The algorithm guarantees the stability of the object pile during packing and the feasibility of the robot motion executing the placement plans. Our formulation for the constraints are simple, but yields surprisingly effective results. Experimental evaluation of the method is conducted with a realistic physical simulator on a dataset of scanned real-world items, demonstrating stable and high-quality packing plans compared with other 3D packing methods. Moreover, a prototype robot hardware and software system that executes the algorithm achieves a

98% success rate with robustness measures. We hope that by highlighting the importance of robot-feasible offline planning and error-correcting measures, the ideas introduced in this thesis can inform key design decisions for future robot warehouse automation technologies. Our contributions are as follows:

1. A set of constraints robot-packable constraints are formulated, and a constructive packing pipeline is proposed to solve these constraints. The pipeline is able to pack geometrically complex, non-convex objects while satisfying stability and robot packability constraints. Experimental evaluation of the algorithm conducted with a realistic physical simulator demonstrates stable and high-quality packing plans compared with other 3D packing methods. This work appeared previously as Wang et al. [WH19b], published in ICRA 2019 with co-author Kris Hauser. This work is presented in Chapter 2.
2. Most existing packing algorithms, including our ICRA 2019 work, assumes a controllable packing sequence. However, in many cases, a controllable arrival sequence may not be practical in a warehouse setting. Therefore, building on top of the ICRA 2019 work, we formulate two variants of packing problems in which the set of items is known, but the arrival order is unknown. The goal is to certify that the items can be packed in a given container, and/or to optimize the size or cost of a container so that that the items are guaranteed to be packable, regardless of arrival order. This work appeared previously as Wang et al. [WH19c], published in RSS 2019 as best paper award nominee with co-author Kris Hauser. This work is presented in Chapter 3.
3. Since the proposed algorithm assumes known 3D models of objects for planning and collision detection. We propose a technique for convenient 3D object model acquisition. In which an object reoriented in front of a video camera with multiple grasps and regrasps. Experiments show that our method results in high quality reconstructed models. Moreover, testing with a novice user on a set of 200 objects

demonstrates relatively rapid construction of complete 3D object models. This work appeared previously as Wang et al. [WH19a], published in ICRA 2019 with co-author Kris Hauser. This work is presented in Chapter 4.

4. Finally, We evaluate the proposed planner under real-world uncertainties such as vision, grasping, and modeling errors. We build hardware and software testbed that is fairly representative of current state-of-the-art sensing, perception, and planning for warehouse manipulation. A systematic evaluation of the testbed is then performed to study the sources of error and models their magnitude. Exhaustive experiments are conducted in Monte Carlo simulation and on the physical testbed to examine the feasibility of the packing placements under open-loop baseline conditions as well as two strategies for improving the robustness of robotic packing. Empirical results demonstrate a success rate of up to 98% can be achieved on a physical robot when using robustness measures despite cascading real-world uncertainties. This work is presented in Chapter 5.

In Chapter 6, we present our design of a robot named SwitchIt that is small, relatively inexpensive, easily mounted on a mobile robot, and actuates buttons reliably. While this is not directly related to robot bin packing, it has many interesting insights on designing an integrated robot system. This work appeared previously as Wang et al. [?], published in ICRA 2018 with co-author Kris Hauser.

Finally, we conclude by summarizing our contributions and discussing open challenges for future work in Chapter 7.

Chapter 2

Offline Bin Packing under Stable and Robot-feasible Constraints

This chapter proposes a formulation of the offline packing problem that is tailored to the automated warehousing domain. Besides minimizing waste space inside a container, the problem requires stability of the object pile during packing and the feasibility of the robot motion executing the placement plans. To address this problem, a set of constraints are formulated, and a constructive packing pipeline is proposed to solve for these constraints. The pipeline is able to pack geometrically complex, non-convex objects while satisfying stability and robot packability constraints. In particular, a new 3D positioning heuristic called Heightmap-Minimization heuristic is proposed, and heightmaps are used to speed up the search. Experimental evaluation of the method is conducted with a realistic physical simulator on a dataset of scanned real-world items, demonstrating stable and high-quality packing plans compared with other 3D packing methods.¹

¹This chapter is reproduced from Fan Wang and Kris Hauser, “Stable Bin Packing of Non-convex 3D Objects with a Robot Manipulator,” in 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 2019, pp. 8698-8704. This work is funded by Amazon Research Award.

2.1 Introduction

Problems that involve the placement of objects within a container or a set of containers are generally referred to as cutting and packing problems. Most existing packing algorithms apply to idealized scenarios, such as rectilinear objects and floating objects not subject to the force of gravity. To perform automatic packing in warehouses using a pre-computed packing plan, several real-world issues need to be addressed, such as stability under force of gravity, and kinematics and clearance issues for the robot.

For a packing plan to be feasible with a robot manipulator, a comprehensive set of constraints need to be formulated.

In addition to the two standard packing constraints:

1. *Noninterference*. Each object is collision free,
2. *Containment*. All objects are placed within the internal space of the container,

we introduce the following constraints necessary for a robot-packable plan:

3. *Stability*. Each object is stable against previously packed objects and the bin itself, and
4. *Manipulation feasibility*. A feasible robot motion exists to load the object into the target placement. The robot must obey kinematic constraints, grasp constraints, and collision constraints during this motion.

In the following sections, we refer to constraints 1 and 2 as the *non-overlap* constraints, and constraints 1-4 as all constraints, or the robot-packable constraints.

While the application of robot-packable constraints is independent of the particular packing problem addressed, this chapter focuses on the problem of offline packing of 3D irregular shapes into a single container. To solve this problem under robot-packable constraints, we present the following main contributions:

1. A polynomial time constructive algorithm to implement a resolution-complete search amongst feasible object placements, under *robot-packable constraints*.
2. A 3D positioning heuristic named Heightmap-Minimization (HM) that minimizes the volume increase of the object pile from the loading direction.
3. A fast prioritized search scheme that first searches for robot-packable placement in a three-dimensional space that likely contains a solution, and falls back to search in a five-dimensional space.

Our algorithm and others in comparison are tested in a realistic physics simulator, by packing large quantities of itemsets using highly complex, real-world object scannings. With item sizes of 3-5 objects (e.g., a common Amazon order size), the success rate is 99.9% for finding and executing packing plans using small Amazon order boxes. Large number of items are also packed in stress tests, in these tests, 80% of the placement plans were successfully executed in the physics simulator, which is significantly better than the 17% success rate from a standard packing solver under the same testing condition. Empirical results also show that the new Heightmap-Minimization heuristic finds more placements than existing heuristics.

2.2 Problem Definition

We address the problem of offline packing of 3D irregular shapes into a single container while ensuring the stability of each packed item and feasibility of the placement with a robot gripper.

Specifically, for a set N geometries $\mathcal{G}_1, \dots, \mathcal{G}_N$ where $\mathcal{G}_i \subset \mathcal{R}^3$, sitting at initial poses $\mathcal{P} \setminus \sqcup = (P_{int1}, \dots, P_{intN})$. Let \mathcal{C} denote the free space volume of the container and $\partial\mathcal{C}$ as the boundary of the free space. Let $T_i \cdot \mathcal{G}_i$ denote the space occupied by item i when the geometry is at pose P_i . The problem is to find a placement sequence $S = (s_1, \dots, s_N)$ of $\{1, \dots, N\}$ and poses $\mathcal{P} = (P_1, \dots, P_N)$ such that each placement satisfies robot-packable

constraint with geometries placed prior:

$$(T_i \cdot \mathcal{G}_i) \cap (P_j \cdot \mathcal{G}_j) = \emptyset, \forall i, j \in \{1, \dots, N\}, i \neq j \quad (2.1)$$

$$T_i \cdot \mathcal{G}_i \subseteq \mathcal{C}, \forall i \in \{1, \dots, N\} \quad (2.2)$$

and for each $k = 1, \dots, N$, stability constraints:

$$isStable(P_{s_k} \cdot \mathcal{G}_{s_k}, \mathcal{C}, T_{s_1} \cdot \mathcal{G}_{s_1}, \dots, P_{s_{k-1}} \cdot \mathcal{G}_{s_{k-1}}) \quad (2.3)$$

and manipulation feasibility constraints:

$$isManipFeasible(P_{s_k} \cdot \mathcal{G}_{s_k}, P_{s_1} \cdot \mathcal{G}_{s_1}, \dots, P_{s_{k-1}} \cdot \mathcal{G}_{s_{k-1}}) \quad (2.4)$$

It is important to note that both stability and manipulation feasibility constraints must be satisfied for *every intermediate arrangement* of objects, not just the final arrangement.

2.2.1 Stability checking

Stability is defined as the condition in which all placed items are in static equilibrium under gravity and frictional contact forces. We model the stack using point contacts with a Coulomb friction model with a known coefficient of static friction. Let the set of contact points be denoted as c_1, \dots, c_K , which have normals n_1, \dots, n_N , and friction coefficients μ_1, \dots, μ_K . For each contact c_k , let the two bodies in contact be denoted A_k and B_k . Let f_1, \dots, f_K denote the contact forces, with the convention that f_k is applied to B_k and the negative is applied to A_k . We also define m_i as the mass of object i , and cm_i as its COM. We take the convention that the container has infinite mass.

The object pile is in static equilibrium if there are a set of forces that satisfy the following conditions.

Force balance: $\forall i = 1, \dots, N$,

$$-\sum_{k|i=A_k} f_k + \sum_{k|i=B_k} f_k + m_i g = 0. \quad (2.5)$$

Torque balance: $\forall i = 1, \dots, N,$

$$- \sum_{k | i=A_k} (cm_i - c_k) \times f_k + \sum_{k | i=B_k} -(cm_i - c_k) \times f_k = 0.$$

Force validity: $\forall k = 1, \dots, K,$

$$f_k \cdot n_k > 0, \tag{2.6}$$

$$\|f_k^\perp\| \leq \mu_k(f_k \cdot n_k). \tag{2.7}$$

where $f_k^\perp = f_k - n_k(f_k \cdot n_k)$ is the tangential component (i.e., frictional force) of f_k .

For a given arrangement of objects, an approximate set of contact points is obtained with the slightly scaled geometries in placement. A pyramidal approximation for the friction cone is used, and the conditions above are formulated as a linear programming problem over f_1, \dots, f_N , solved using the convex programming solver CVXPY [DB16]. If no such forces can be found, the arrangement is considered unstable.

2.2.2 Manipulation feasibility

This constraint checks feasibility of a packing pose when executed by a robot manipulator. This requires that the object be graspable from its initial pose and can be packed in the desired pose via a continuous motion, without colliding with environmental obstacles.

In our system, we limit ourselves to the existence of a feasible top-down placement trajectory within the grasp constraints, as robots performing pick and place (e.g., box packing) commonly use vertical motion [dBKM⁺05]. We also assume the existence of a grasp generator that produces some number of candidate end effector (EE) transforms, specified relative to an object’s geometry that may be used to grasp the object. The pseudo-code for this procedure is given in Alg. 1.

2.3 Pipeline for Robot-packable Planning

We develop a constructive packing pipeline to solve for the set of robot-packable constraints proposed. Our algorithm accepts an itemset, a container dimension, a constructive posi-

Algorithm 1: isManipFeasible

input : Desired placed geometry $T \cdot \mathcal{G}$ and a set of grasp candidates $\{T_1^{\mathcal{G}}, \dots, T_n^{\mathcal{G}}\}$

```
1 for  $T^{\mathcal{G}} \in \{T_1^{\mathcal{G}}, \dots, T_n^{\mathcal{G}}\}$  do
2   Compute top-down EE path  $\mathcal{P}_{ee}$  interpolating from an elevated pose to a
   final pose  $T \cdot T^{\mathcal{G}}$ ;
3   for  $P_{ee} \in \mathcal{P}_{ee}$  do
4     if  $\neg(\text{IKSolvable}(P_{ee}) \wedge \text{inJointLimits}(P_{ee}) \wedge \text{collisionFree}(P_{ee}))$  then
       Continue with Line1;
5   end
6   return True
7 end
8 return False
```

tioning heuristic, and/or a packing sequence, to produce packing plans. The pipeline packs each item to its optimized feasible pose in sequential order, without backtracking.

Our pipeline primarily consists of 4 components, namely:

1. Placement sequence
2. Generate ranked transforms
3. Stability check
4. Manipulation feasibility check

The pipeline starts with a sequencing heuristic to sort all items in a tentative placement ordering and allocates them individually into the container in this sequence. For each object at the time of the allocation, a set of candidate transforms satisfying robot-packable constraints are generated and ranked based on the positioning heuristic. Constraint checks are performed in order until a transform satisfying all required constraints is returned.

2.3.1 Placement sequence

The placement sequence can be user-specified or generated by non-increasing bounding box volume rule. The generated sequence is subject to adjustment if a solution cannot be found in the specified ordering.

2.3.2 Generating ranked transforms

For a given item, a positioning heuristic (e.g., placement rule) identifies a free pose inside the container that is most preferred according to a specific criterion. Our pipeline accepts arbitrary positioning heuristics, but instead of applying the heuristic to obtain one optimal placement for each item, we use the score formulated from the positioning heuristic to rank candidate placements.

The candidate placements are obtained with a prioritized search among a discretized set of object poses. Instead of searching in the 6D space of $SE(3)$, our algorithm first performs a grid search in a 3D space that likely contains robot-packable solutions. In the 3D search, the rolls and pitches of G are restricted to be a set of *planar-stable* orientations, which are a set of stable resting orientations of G on a planar surface, computed using the method of Goldberg et al. [GMZ⁺99]. This speeds up the search for the common case of packing on the first layer and on horizontal supports. If no feasible solutions exist in the 3D space, the algorithm falls back to search in 5D, in which a grid search is performed for rolls and pitches as well.

The 3D search for collision-free placements of one object, given a set of rolls and pitches, is shown in Alg. 2. A grid search is performed for yaw, X , and Y at a given resolution, and the height Z of the placement is analytically determined as the lowest free placement. 2D heightmaps are used to accelerate the computation of Z to an efficient 2D matrix manipulation. Three heightmaps are computed: 1) a top-down heightmap H_c of the container and placed objects, 2) a top-down heightmap H_t of the object to be placed, and 3) a bottom-up heightmap H_b of the object to be placed. H_t and H_b are measured relative to the lower left corner of the orientated object. Raycasting is used to build these heightmaps, and rays

that do not intersect with the object geometry are given height 0 in H_t and ∞ in H_b . The container heightmap is obtained once at the beginning of object placement search, and an object heightmap is computed once for each distinct searched orientation.

Given an object orientation and X, Y location, we calculate the lowest collision-free Z as follows:

$$Z = \max_{i=0}^{w-1} \max_{j=0}^{h-1} (H_c[x+i, y+j] - H_b[i, j]) \quad (2.8)$$

where (x, y) are the pixel coordinates of X, Y , and (w, h) to be the dimensions of H_t .

Algorithm 2: 3DGridSearch

input : Geometry \mathcal{G} , container \mathcal{C} , rolls and pitches O

output: All collision-free candidate transforms

```

1 for  $(\phi, \psi) \in O$  do
2   for  $\theta \in \{0, \Delta r, 2\Delta r, \dots, 2\pi - \Delta r\}$  do
3     Let  $R \leftarrow R_z(\theta)R_y(\phi)R_x(\psi)$ ;
4     Discretize legal horizontal translations of  $R \cdot \mathcal{G}$  into grid
        $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$ ;
5     for  $(X, Y)$  in  $\{(X_1, Y_1), \dots, (X_n, Y_n)\}$  do
6       Find the lowest collision free placement  $Z$  at translation  $X, Y$ ;
7        $T \leftarrow (R, (X, Y, Z))$ ;
8       if  $T \cdot \mathcal{G}$  lies within  $\mathcal{C}$  then Add  $T$  to  $\mathcal{T}$ ;
9     end
10  end
11 end
12 return  $\mathcal{T}$ 

```

Once all collision-free candidate transforms are obtained, they are scored by a scoring function formulated from a positioning heuristic. For example, the Deepest-Bottom-Left-

First heuristic can be formulated as the score:

$$Z + c \cdot (X + Y) \tag{2.9}$$

where c is a small constant.

The candidates are then ranked by score (lower is better). If only robot-packable constraints are required, the placement candidate with the lowest score is returned. If additional constraints are specified, the ranked candidates will be checked for the additional constraints until a candidate satisfying all constraints is returned.

After a new object has been placed, we update the heightmap of the container H_c . This subroutine is also used in our heightmap minimization heuristic. Given a pose X, Y, Z of the object to be packed, and the top heightmap H_t at the given orientation, we calculate an updated heightmap H'_c adding the placed object as follows:

For all $i = 0, \dots, w - 1, j = 0, \dots, h - 1$, we let:

$$H'_c[x + i, y + j] = \max(H_t[i, j] + Z, H_c[x + i, y + j]) \tag{2.10}$$

if $H_t[i, j] \neq 0$, and otherwise

$$H'_c[x + i, y + j] = H_c[x + i, y + j]. \tag{2.11}$$

2.3.3 Pipeline summary and fall back procedures

A packing attempt for a single item is summarized in Alg. 3, and the overall pipeline for packing multiple objects is given in Alg. 4. Given a heuristic packing sequence, it calls Alg. 3 for each item with the set of planar-stable rolls and pitches. This first stage finds placements for most objects in typical cases. For the remaining unpacked items U , the algorithm activates the *fallback procedure*. The fallback procedure examines each unpacked item and attempts to perturb the planar stable orientations by iterating over rolls and pitches until a solution is found, and if no solution is found the algorithm terminates with failure.

Algorithm 3: packOneItem

input : item geometry \mathcal{G} , container \mathcal{C} , pitches and yaws O , sequence of the packed items $\{s_1, \dots, s_i\}$, transforms of the packed items $\{P_1, \dots, P_i\}$

output: Transform T or **None**

- 1 $\mathcal{T} \leftarrow \text{3DGridSearch}(\mathcal{G}, \mathcal{C}, O)$;
- 2 Score each T in \mathcal{T} based on heuristic used;
- 3 **for** up to N lowest values of T in \mathcal{T} **do**
- 4 **if** $\neg \text{isStable}(T \cdot \mathcal{G}, \mathcal{C}, P_1 \cdot \mathcal{G}_{s_1}, \dots, P_i \cdot \mathcal{G}_{s_i})$ **then continue**;
- 5 Compute grasp poses $T_1^{\mathcal{G}}, \dots, T_n^{\mathcal{G}}$ compatible with T ;
- 6 **if** $\text{isManipFeasible}(T \cdot \mathcal{G}, \{T_1^{\mathcal{G}}, \dots, T_n^{\mathcal{G}}\})$ **then return** T ;
- 7 **end**
- 8 **return** *None*

2.4 Heightmap-Minimization Heuristic

The performance and solution quality of a multi-dimensional packing problem is highly susceptible to the item-positioning rule [LMV04]. However, existing positioning heuristics for 3D packing are scarce and are commonly adapted directly from 2D packing, and therefore result in poor space utilization in the 3D container [LMV02, CPT08]. To address these shortcomings, we propose a novel positioning heuristic called the Heightmap-Minimization (HM) heuristic, which favors item placements that result in the smallest occupied volume in the container, as observed from the loading direction.

Specifically, HM scores a placement as follows. Given the candidate transform $T = (\text{roll}, \text{pitch}, \text{yaw}, X, Y, Z)$, compute a tentative container heightmap H'_c using the update routine described in Sec. 12. Suppose its shape is (w, h) . The score for the placement using the HM heuristic is:

$$c \cdot (X + Y) + \sum_{i=0}^{w-1} \sum_{j=0}^{h-1} H'_c[i, j] \quad (2.12)$$

Algorithm 4: Robot-feasible packing with fall back procedures

input : Item geometries $\mathcal{G}_1, \dots, \mathcal{G}_N$, container C , initial packing sequence $\{s0_1, \dots, s0_N\}$

output: Transforms \mathcal{T} and sequence S , or **None**

- 1 Initialize $\mathcal{T}, S, U, \mathcal{O}$ to empty lists;
- 2 **for** $\mathcal{G}_i \in \{\mathcal{G}_1, \dots, \mathcal{G}_N\}$ **do**
 - 3 Get planar-stable rolls and pitches for \mathcal{G}_i with the top n highest quasi-static probabilities $O_i = \{(\phi_1, \psi_1), \dots, (\phi_n, \psi_n)\}$;
 - 4 Add O_i to \mathcal{O} ;
- 5 **end**
- 6 **for** $(s0_i \in \{s0_1, \dots, s0_N\})$ **do**
 - 7 $T = \text{packOneItem}(G_{s0_i}, C, O_{s0_i}, S, \mathcal{T})$;
 - 8 **if** T **then** Add T to \mathcal{T} , Add $s0_i$ to S ;
 - 9 **else** Add $s0_i$ to U ;
- 10 **end**
- 11 **for** $u_i \in U$ **do**
 - 12 Let $\{(\phi_1, \psi_1), \dots, (\phi_n, \psi_n)\}$ be the planar-stable orientations in O_{u_i} ;
 - 13 **for** $t_r \in \{0, \Delta r, 2\Delta r, \dots, 2\pi - \Delta r\}$ **do**
 - 14 **for** $t_p \in \{0, \Delta r, 2\Delta r, \dots, 2\pi - \Delta r\}$ **do**
 - 15 $O^t = \{(\phi_1 + t_r, \psi_1 + t_p), \dots, (\phi_n + t_r, \psi_n + t_p)\}$;
 - 16 $T = \text{packOneItem}(G_{u_i}, C, O^t, S, \mathcal{T})$;
 - 17 **if** T **then**
 - 18 Add T to \mathcal{T} ; Add u_i to S ;
 - 19 **continue** with Line 11
 - 20 **end**
 - 21 **end**
 - 22 **return** *None*
 - 23 **end**
 - 24 **return** (\mathcal{T}, S)




Figure 2.1: Example packing placements obtained by HM and DBLF. HM finds more compact and stable packing compared to DBLF.

where c is a small constant.

HM favors positions and orientations that result in good space utilization as it minimizes wasted space and holes that cannot be filled. HM also favors stable placements since the bottom of the object is encouraged to match the shape of the supporting terrain (Fig. 2.1).

2.5 Experiment

We tested our algorithm on different item sets and validated plan feasibility in a physics simulator. Objects were drawn at random from a set of 94 real-world object meshes from the YCB [CSB⁺17] and the APC 2015 object set [Rut15]. On average each mesh contains 10,243 vertices. Experiments are conducted on Amazon Web Services instance type m5.12xlarge. All computation times are measured on a single thread. Parameters used in the experiment are: heuristic constant $c = 1$; heightmap resolution 0.002m; step size in both X and Y 0.01m; $\Delta r = \pi/4$ in range $[0, \pi)$; friction coefficient $\mu = 0.7$. Contact points are obtained using the exact geometry with a scale factor of 1.03. The top 4 planar-stable rolls and pitches with the highest quasi-static probabilities are used, and candidate number $N = 100$.




Figure 2.2: (a) Grasp poses generated satisfying vacuum graspability constraints. (b) Compatible gripper poses with candidate object orientation are checked for clearance with the container and the object pile. Collision-free grasps are shown in green and colliding grasps are colored in transparent grey.

2.5.1 Robot manipulation feasibility with a vacuum gripper

The robot model used to verify robot feasibility constraints is a Staubli TX90 robot, equipped with a cylindrical vacuum gripper of 30 cm length and 2 cm diameter. The graspability constraints ask the vacuum gripper to grasp within a radius $r = 2$ cm in the horizontal plane to the object's center of mass when the object is sitting in flat orientations. The areas under the gripper should be solid planar areas (80% of the surface points directly below the tool are within 0.3mm to the estimated plane) for the vacuum opening to grasp normal to the surface, the resulting gripper axis needs to be within a tilting angle $\theta = \pi/4$ to the Z-axis.




Figure 2.3: Distribution of the solution containers found under different level of constraints. The x axis is the 5 container dimensions tested.

2.5.2 Small Order Packing

Simulating problem settings in a typical warehouse, we performed a small order packing test. Per communication with personnel at Amazon, 3-5 items are a standard order size. We generated 1000 random itemsets consisting of 3-5 models. The itemsets are verified with various testing methods to fit in at least one of the five containers used in the Amazon Robotics Challenge 2017 [ARC] under robot-packable constraints, and the smallest feasible container for each itemset is recorded.

We test our pipeline with HM heuristic and all constraints. Our algorithm seeks a feasible solution using the smallest container first, and if fails, the process is repeated on the second smallest container and so on until either a solution is found or all available containers are exhausted.

The success rate is 99.9% averaging 9.54s per order. The only failure occurs when a large object needs to be tilted sideways to fit within a tight space, and no feasible vacuum grasp exists in the specified orientation. Fig. 2.3 compares the smallest bin statistics using all tested methods and ours. It appears that HM packing, despite adding all constraints, enlarges the container needed only marginally.



Figure 2.4: Examples of packing plans for itemsets of size 10.

2.5.3 Comparisons on Large Itemsets

Next, we perform stress tests on itemsets of size 10. A tall container of size $32 \times 32 \times 30$ cm is chosen. 1000 itemsets of size 10 are generated and verified with all tested methods to have a non-overlap packing within the chosen container. Since the tilted gripper is likely to collide with the tall container chosen, we assume the gripper can grasp object of any orientation at the center of object’s top surface, with the gripper axis vertically aligned to the Z axis.

We compare our HM heuristic against the DBLF and MTA heuristics [WGC⁺10], as well as an implementation of a guided local search (GLS) method as described by Egeblad et al. [Ege09]. The fast intersection area theorem in Egeblad’s paper was not implemented. Therefore, for the fairness of the comparison, GLS was run with 5 random restarts, and each restart was terminated after 300s if a solution could not be obtained. GLS is also not tested for all constraints, as implementing robot-packable constraints in GLS methods is very challenging. Table 5.2 reports the percentage of solutions found and the average computation time.

Empirically, HM finds more solutions than any other method in comparison. With robot-packable constraint, HM finds 99.9% of all feasible solutions, leading the 2nd place DBLF heuristic by 1.5%, while MTA and GLS are not as competitive. After adding all constraints, each technique drops in success rate by a few percents, but HM still leads the other methods. The mean running time of HM is also 30% shorter, indicating that the

Table 2.1: Comparing planning techniques on 10-item orders with and without robot-packable constraints

	HM	DBLF [WGC+10]	MTA [WGC+10]	GLS [ENO07]
Success, non-overlap (%)	99.9	98.4	88.9	78.9
Time, non-overlap (s)	15.7	14.2	14.1	502
Success, all constraints (%)	97.1	96.3	86.3	—
Time, all constraints (s)	34.9	50.1	95.4	—

Table 2.2: Impact of Δr and candidate number N on the results

Rotation granularity Δr	$\pi/4$	$\pi/4$	$\pi/8$
Number of candidates N	100	500	500
Success, all constraints (%)	97.1	97.5	98.7
Time, all constraints (s)	34.9	70.05	89.40

highest ranked placements are more likely to be stable than the other heuristics.

In addition, only 3.2% (320 out of 10,000) of the items are packed with the fallback procedure, indicating the 3D space searched is indeed highly likely to contain robot-packable solutions. The fallback procedure is nonetheless important, as, with no fallback procedure, the success rate with all constraints drops from 97.1% to 72.4%.

With finer rotation granularity Δr and more candidates to check against robot-packable constraints, the success rate can be further improved at the cost of increased computation time (Table 2.2).

Table 2.3: Execution success rates in simulation, 10-item orders

	Success (%)	Drop (cm)	Horiz. Shift (cm)
Non-overlap constraint	17.11	1.95	1.29
All constraint	79.1	1.36	0.50

2.5.4 Executing Packing Plans in Simulation

Finally, we test the open-loop execution feasibility of packing plans in the Klamp't robot physics simulator [kla]. In the simulation, the robot places one item after another using a top-down loading direction. The plan is considered a success if: 1) All items placed to the planned transforms without the robot and the objects colliding with items placed prior, and 2) all items contained within the container when placement is complete.

The robot used in the simulation is the TX90 robot model with the vacuum gripper described. We make the same assumption for the gripper as in the 10-item packing case. The robot places all items 1cm elevated from their planned transform; therefore there is an expected 1 cm drop. We allow 20s for the items to settle before the next item is placed.

In the 3-5 items case, 100% of plans are executed successfully according to our success criteria. In the 10 item case, 768 out of 971 ($\approx 80\%$) of robot-packable plans obtained with HM heuristic are executed successfully. This is significantly higher than the 17% success rate with robot-packable constraints. Further shown in Table 2.3, items packed with all constraints undergo smaller displacements during packing execution, indicating increased stability.

The 20% failure cases are caused by an object falling out of its desired placement, which prevents subsequent items from being packed. The stability checker may be too optimistic, especially for intrinsically unstable objects like balls. Moreover, the impact of dropping an object could shift supporting objects.

Chapter 3

Robot Packing with Known Items and Nondeterministic Arrival Order

Building on top of the offline algorithm introduced in Chapter 2. This chapter formulates two variants of packing problems in which the set of items is known but the arrival order is unknown. The goal is to certify that the items can be packed in a given container, and/or to optimize the size or cost of a container so that that the items are guaranteed to be packable, regardless of arrival order. The Nondeterministically ordered packing (NDOP) variant asks to generate a certificate that a packing plan exists for every ordering of items. Quasi-online packing (QOP) asks to generate a partially-observable packing policy that chooses the item location as each subsequent item is revealed. Theoretical analysis demonstrates that even the simple subproblem of verifying feasibility of a packing policy is NP-complete. Despite this worst-case complexity, practical solvers for both NDOP and QOP are developed, and experiments demonstrate their application to packing irregular 3D shapes with manipulator loading constraints.¹


¹This chapter is reproduced from Fan Wang and Kris Hauser, “Robot Packing with Known Items and Nondeterministic Arrival Order” in *Robotics: Science and Systems(RSS)*, June 2019. The journal version of the paper is currently in submission to *IEEE Transactions on Robotics (T-RO)*. This work is supported by an Amazon Research Award.

3.1 Introduction

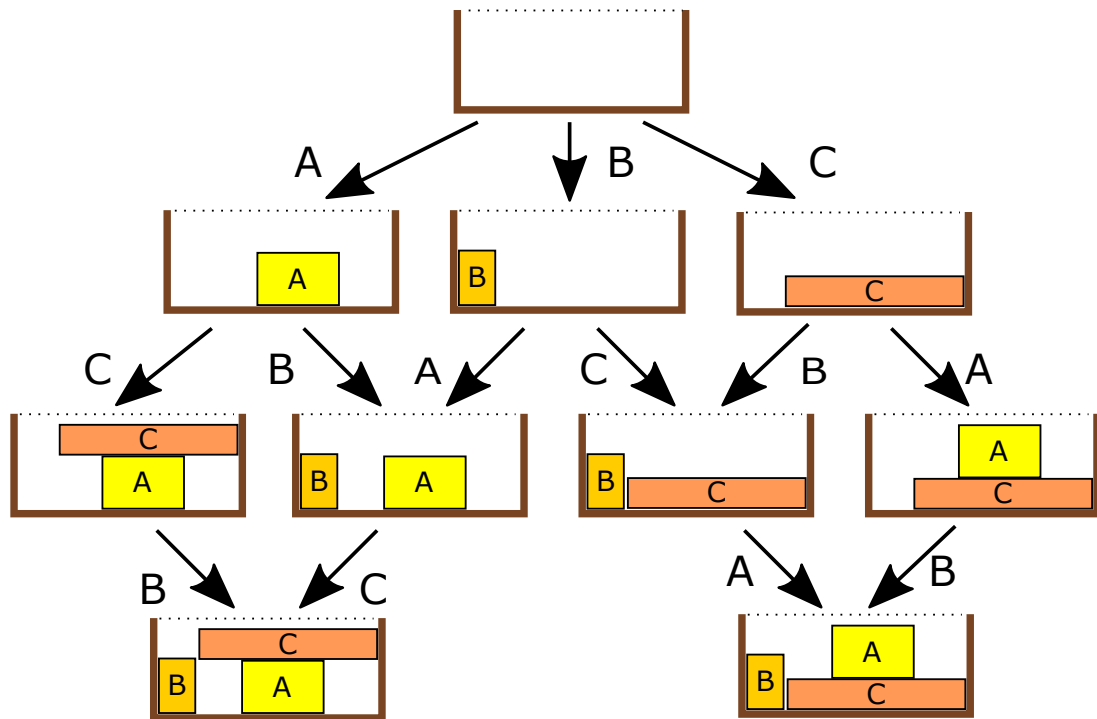
A large variety of packing problems have been studied, including the bin and strip packing problem, knapsack problem, container loading problem, nesting problem, and others. In an *offline* setting, the items and container(s) are known, and a plan can place the items in arbitrary order [MPV00]. In an *online* setting, the items are not known a priori and need to be placed as they arrive [Sei02]. We consider a *robot packing* setting which addresses packing problems with the additional constraints that items must be loaded with a collision-free robot path, and that intermediate piles of items must be stable against gravity.

This chapter introduces two *nondeterministic* formulations of robot packing problems that lie between the offline and online settings. These formulations are practical for automated warehouses where the ultimate item set (e.g., shopping cart) is known, but some distinct, uncontrollable component of the packing system controls the item arrival order. For example, in Amazon’s automated fulfillment centers, shelving units containing individual items are carried by thousands of mobile robots to several picking stations, and the order in which shelves arrive at a given station is controlled by a complex algorithm that is tuned to maximize delivery throughput for shelving units. In cases where item deliveries are human-controlled, it may be even less practical for an algorithm to dictate the arrival order. Hence, to guarantee that the items can fit in a given container, a packing planner should *certify* the validity of a plan under *all possible arrival orders*. In the NDOP variant, the feasibility of the container is verified under all nondeterministic orders, but the arrival sequence is revealed before packing is executed. In the QOP variant, each object must be packed before the next item is revealed (Fig. 3.1).

We present a practical framework for solving NDOP and QOP problems that uses a combination of an offline planner and a packing policy verifier. A packing policy is represented by a set of possible packing plans, each of which consists of a set of packing locations and a directed acyclic graph (DAG) of their dependencies. The verifier will verify or disprove the feasibility of a policy under all permutations of arrival orders. We present a verification algorithm that uses pruning techniques, and in practice can check feasibility



(a) Nondeterministically ordered packing (NDOP)



(b) Quasi-online packing (QOP)

Figure 3.1: Feasible solutions for a 2D, 3-item instance of (a) NDOP and (b) QOP. All $3! = 6$ possible arrival orders are collision-free, loadable from top-down, and yield intermediate piles that are stable under gravity. In QOP, an item is never moved after it is placed.

quickly even for a large number of objects and packing plans. However, in some cases exponential behavior is observed. We prove that the worst-case solution complexity of NDOP and QOP is $O(n!)$ and even feasibility verification for a polynomial-sized NDOP policy is NP-complete, via reduction from SAT.

Nevertheless, the solver is practical for small numbers of items, and even using an incomplete offline planner, it guarantees that a solution, when found, is feasible for all object orderings. Several packing heuristics are also introduced to improve scalability of the approach, and experiments demonstrate that our approach can be realistically applied to irregular 3D shapes with item sets of size up to 10.

3.2 Related work

Packing algorithms have been studied extensively both for their theoretical interest and practical applications in shipping, manufacturing, and 3D printing. The vast majority of work considers rectilinear objects. State-of-the-art exact algorithms for the offline 2D and 3D bin packing problem use branch-and-bound approaches [MV98, MPV00]. Because exact methods have worst-case exponential complexity, heuristic methods and metaheuristic approaches have been developed, such as the Bottom-Left [BCR80] and Best-Fit-Decreasing heuristics [JDU⁺74]. Heuristics are the only practical methods available for irregular shape packing (a.k.a. nesting [ENO07]), since the freedom to rotate leads to a continuously infinite search space. Metaheuristic optimization methods [FPZ03, kam88] simultaneously optimize the placements of all items, and constructive heuristics incrementally place items according to some scoring function [LLCY15, WGC⁺10].

In the offline setting, the item set and packing order can be controlled. Most classical versions do not formulate interdependence between items (i.e., items appear and then “float” in their planned locations), which means ordering is irrelevant. This is a reasonable assumption for some scenarios like sheet metal cutting, but one should consider additional constraints when packing containers in practice. Prior work has enforced clearance of

boxes along axis-aligned loading directions in 3D bin-packing by ensuring no previously-placed item lies along an extruded prism along at least one face of the box [dBKM⁺05]. Recent work has also formulate collision-checking between the loading mechanism (human hand, forklift, robot hand, etc.) and already-placed items [WH19b]. Stability constraints [Ege09, LLCY15, WH19b] also impose dependency on the packing order. Due to these dependencies, if the item arrival order does not match the planned order, the plan might not be successfully executed.

In the online setting, an arbitrary item is presented to the algorithm, which then chooses a packing location [Sei02]. Neither the item set nor the packing order are controllable, and often the item geometry can also be arbitrary as well. There are no guarantees that a given container can be packed, so the typical formulation casts the problem as an optimization of the number of containers or the container height. Competitive ratios are known for various online algorithms in the 1D and 2D rectilinear bin packing settings [HCT⁺11, Sei02], but to our knowledge, no results are known for irregular shapes. In contrast, NDOP and QOP seek guaranteed packing in a single container when the item set is known, which is more appropriate for fulfillment applications. Our NDOP and QOP solvers can handle irregular shapes, and we prove that when they successfully return a policy, the answer is correct.

3.3 Problem formulation

Let $\mathcal{I} = \{v_1, \dots, v_n\}$ be a set of n items. Item v_i has some geometry $A_i \subset \mathbb{R}^d$, and we wish to pack all items into a container volume $C \subset \mathbb{R}^d$. Here $d = 2$ or 3 is the dimension of the workspace. The offline packing problem is to compute a *feasible packing plan* given A_1, \dots, A_n and C . Such a plan is defined as follows:

Definition 1. A packing plan P consists of an ordering $\sigma_{1:n} = (\sigma_1, \dots, \sigma_n)$ and a tuple of transforms $T_{1:n} = (T_1, \dots, T_n)$, in which $\sigma_j \in \{1, \dots, n\}$ specifies that v_{σ_j} is the j 'th item to be placed, and $T_i \in SE(d)$ specifies the target location (pose) of v_i .

An ordering must be a permutation on n elements, and is hence an element of the




Figure 3.2: Examples of plans that are infeasible for arrival order ABCD: (a) unstable, (b) items C and D collide with B along the loading direction, (c) and the path for the robot manipulator to grasp and load item D is infeasible. In (d), although ABCD is feasible, the prefix requirement is violated because the sub-plan ABC is unstable.

symmetric group S_n

Definition 2. A packing plan is feasible when it, and all prefix plans, satisfy certain constraints, as shown in Fig. 3.2 and defined in the below section.

The prefix feasibility requirement means that for all $j < n$, the ordering $\sigma_{1:j}$ with the corresponding items in locations $T_{\sigma_1}, \dots, T_{\sigma_j}$ must also satisfy the feasibility constraints. For example, we cannot require two blocks to be placed simultaneously on either ends of a see-saw when stability is violated with only a single block (Fig. 3.2.d).

3.3.1 Constraint formulation

In our formulation, the feasibility of a packing plan requires satisfying the following three constraints. For readability, for the ordering $\sigma_{1:n}$ let us denote the sequence number s_i of

the i 'th item to be the ordinal index in which it appears, i.e., $s_{\sigma_j} = j$ and $\sigma_{s_i} = i$.

Non interference All items do not overlap but can touch (3.1) and all items lie entirely inside the container (3.2):

$$T_i A_i^\circ \cap T_j A_j^\circ = \emptyset \text{ for all } i, j \text{ with } i \neq j, \quad (3.1)$$

$$T_i A_i \subseteq C \text{ for all } i. \quad (3.2)$$

Here \cdot° denotes a set's interior.

Equilibrium To prevent ‘‘floating’’ items and unbalanced stacks, the equilibrium constraint requires that each intermediate packing be stable under gravity and frictional contact (Fig. 3.2.a). An item is allowed to make contact with the container walls and previously placed items. We model these as a set of contact points, and require that there exist feasible forces at each contact point that respect Coulomb friction.

Manipulation feasibility Each item in the packing plan must be loadable by a manipulator without disturbing previously packed items (Figs. 3.2.b and 3.2.c). We consider a robot gripper R and a top-down loading direction. In the packing plan, an item is also given a grasp transform T_G , such that the combined geometry of the i th item and the robot while grasped is $A_i \cup T_G R$. The swept volume of the item and robot while loading is $SV_i = \overline{ab} \oplus T_i(A_i \cup T_G R)$, where $a = (0, 0, 0)$ and $b = (0, 0, h)$, with h some ‘‘safe’’ height greater than the height of the container. This constraint states that, for all items i , the swept volume cannot intersect any previously-placed items (3.3) or the container walls ∂C (3.4):

$$SV_i^\circ \cap T_j A_j = \emptyset \text{ for all } j \text{ s.t. } s_j < s_i, \quad (3.3)$$

$$SV_i^\circ \cap \partial C = \emptyset. \quad (3.4)$$

3.3.2 Nondeterministic problems

Definition 3 (NDOP). *The nondeterministically ordered packing problem asks whether there exists a feasible packing plan for every ordering $\sigma_{1:n} \in S_n$.*

To define QOP, we need to define the concept of a *feasible packing policy* as follows:

Definition 4. *A packing policy is a function π that takes as arguments the identities and locations of previously packed items $(T_{\sigma_1}, \dots, T_{\sigma_{j-1}})$ and the next item σ_j to be packed, and returns the location T_{σ_j} of the next packed item.*

Definition 5. *A packing plan is generated by a packing policy π and an ordering $\sigma_{1:n} \in S_n$ via the recursive application of the policy: $T_{\sigma_1} = \pi((), \sigma_1)$, $T_{\sigma_2} = \pi((T_{\sigma_1}), \sigma_2)$, \dots , $T_{\sigma_n} = \pi((T_{\sigma_1}, \dots, T_{\sigma_{n-1}}), \sigma_n)$.*

Definition 6. *A packing policy is feasible if for all item orders $\sigma_{1:n} \in S_n$, the generated packing plan is feasible.*

Since a packing policy is deterministic, we can also write the policy as a function of the prior order of the objects:

$$\pi((\sigma_1, \dots, \sigma_{j-1}), \sigma_j) \equiv \pi((T_{\sigma_1}, \dots, T_{\sigma_{j-1}}), \sigma_j). \quad (3.5)$$

A policy can also be viewed as a tree with depth n and each node has $n - \ell$ branches on level ℓ . This gives a total of $\sum_{\ell=1}^n n!/\ell! = O(n \cdot n!)$ nodes altogether.

Definition 7 (QOP). *The quasi-online packing problem asks to compute a feasible packing policy.*

The main difference between NDOP and QOP is that with QOP, the items are revealed in sequence, and the location chosen for an item is fixed and may not be changed thereafter. QOP is at least as hard as NDOP, because any solution to QOP is also a solution to NDOP (but the converse does not hold).

3.3.3 Container optimization variants

Above we have stated these packing problems in their decision versions. We also consider container optimization variants, which assume a set of possible containers \mathcal{C} and a cost function $cost : \mathcal{C} \rightarrow \mathbb{R}$, and are stated as follows:

- Offline: Find the container $C \in \mathcal{C}$ with minimum cost that yields a feasible packing plan for item set \mathcal{I} .
- Nondeterministically-ordered: Find the container $C \in \mathcal{C}$ with minimum cost that yields a feasible packing plan for any ordering of item set \mathcal{I} .
- Quasi-online: Find the container $C \in \mathcal{C}$ with minimum cost that yields a feasible packing policy for item set \mathcal{I} .

The container set is typically discrete, such as a set of available boxes, but could also be continuous, such as a varying height. This formulation can express the classical bin-packing problem, where \mathcal{C} contains a container with 1 bin, a container with 2 bins, and so on, and cost measures the number of bins.

NDOP and QOP are adapted rather easily into discrete container optimization algorithms by enumerating containers in order of non-decreasing cost until a successful packing policy is found.

3.4 Method

We make use of an offline robot packing planner [WH19b] with a small amount of modification. The responsibility of the offline planner is to generate a feasible packing plan given the constraints outlined above, while our key contributions are novel methods to invoke the planner and to validate plans under permutations of item orders.

3.4.1 Offline planner

The offline planner is required to accept some number of fixed items and a partial packing sequence for the remaining items. Its interface takes the form

$$P \leftarrow \text{Offline-Pack}(\sigma_{1:j}^{fixed}, P^{prior}, \sigma_{j+1:k}^{next}) \quad (3.6)$$

producing either a feasible plan $P = (\sigma_{1:n}, T_{1:n})$ or “failure.” The inputs $\sigma_{1:j}^{fixed}$ specify that j items of the prior plan P^{prior} should be kept in their previous positions, and $\sigma_{j+1:k}^{next}$ are a sequence of $k - j > 0$ items that should be placed next. The remaining $n - k$ items can be placed in arbitrary order. Specifically, the result must satisfy $\sigma_{1:j} = \sigma_{1:j}^{fixed}$, $\sigma_{j+1:k} = \sigma_{j+1:k}^{next}$, and each fixed transform T_j for $j \in \sigma_{1:k}^{fixed}$ matches the corresponding transform in P^{prior} .

The offline planner used here is a constructive, heuristic method that is easily modified to handle the required changes. In Sec. 3.5 we also consider offline packing heuristics that make the job of generating a nondeterministic plan easier, but these are not strictly necessary.

3.4.2 Compatibility

A naive algorithm to solve NDOP would compute a packing plan for all $n!$ orderings. However, the notion of *plan compatibility* allows us to validate large numbers of orderings for lightly-interdependent plans. For example, if we ignored manipulation feasibility and equilibrium constraints, there is no sequential dependence between any two items, and hence all orderings of items would be feasible under the following policy: *when an item arrives, just place it in its planned location*. We define compatibility as follows:

Definition 8 (Compatible ordering). *A packing plan $P = (\sigma_{1:n}, T_{1:n})$ is compatible with an ordering $\sigma'_{1:n}$ if the reordered plan $P' = (\sigma'_{1:n}, T_{1:n})$ is feasible.*

Hence, we can recast the problem of generating a feasible packing policy as one of generating a set of feasible plans with sufficient coverage as follows:




Figure 3.3: A plan and its dependency graph. C_2 requires D_1 to be present to maintain the stability constraint, because otherwise the imbalanced weight on B would cause tipping. Similarly, D_2 depends on C_1 , and so forth for C_3 and D_3 . This CDG is compatible with orders of the form $AB(C_1D_1)(C_2D_2)(C_3D_3)$ where the (XY) denotes either XY or YX .

Definition 9 (NDOP #2). *Compute a set of feasible plans P_1, \dots, P_m such that for any order $\sigma_{1:n} \in S_n$, there is at least one plan compatible with $\sigma_{1:n}$.*

Our NDOP solver formulates a packing policy as a set of packing plans P_1, \dots, P_m along with their associated *constraint dependency graphs* (CDGs) G_1, \dots, G_m as defined in Sec. 3.4.3. An individual packing plan can be used for the set of orderings that are compatible with its dependency graph. If the union of the m sets of compatible orderings covers S_n , then we are done. If not, we find an incompatible ordering using Alg. 5, and generate a new plan for this ordering.

A QOP solver must address the problem that if any two plans share the same order prefix, each item location in the prefix must be the same. Our algorithm uses the same CDG data structure to calculate compatibility while generating an optimized policy tree.

3.4.3 Constraint dependency graphs

A fundamental data structure that will allow us to verify compatibility is the constraint dependency graph (CDG). This structure (Fig. 3.3) explicitly models the dependencies between items, so that compatibility can be quickly verified.

Definition 10 (CDG). *The CDG of a feasible plan P is a graph on vertices \mathcal{I} that has an edge (u, v) if some feasibility constraint requires item u to be placed before item v .*

We can see that a CDG is a directed acyclic graph (DAG), because if there were a cycle in the graph, by transitivity an item on the cycle would need to be placed before itself. Moreover, a CDG can be replaced by its transitive reduction with no loss in compatibility information.

To construct a CDG $G = (\mathcal{I}, E)$ of a plan $P = (\sigma_{1:n}, T_{1:n})$, we do so in incremental fashion by testing all pairwise constraints. Observe that there is no edge $(\sigma_j, \sigma_i) \in E$ for $i < j$, and we need not add edges (u, σ_i) for any ancestors of σ_i already in the CDG. For each index i in increasing order, we check all $u \in \sigma_{1:i-1}$ for a dependency in reverse packing order. First, if u is an ancestor of σ_i , it is skipped because σ_i is already dependent on u . Next, the manipulation feasibility constraint of u is checked against σ_i . If so, we add an edge (u, σ_i) . If not, we proceed to check equilibrium of the partial stack that includes $\sigma_{1:i}$ but omits u and all descendants of u . If there is no equilibrium solution, we add an edge (u, σ_i) (see Fig. 3.3). It should be noted that there exist scenarios that are stable if a single predecessor item is removed, but unstable if multiple predecessors are removed. These examples, however, are convoluted “multiple see-saw” constructions, and would be highly unlikely to be generated by an offline packing planner.

An ordering is compatible with a plan P iff it does not violate any dependency in P 's CDG. In other words, a feasible packing plan P with a dependency-free CDG $G = (\mathcal{I}, \emptyset)$ is compatible with all orderings. More precisely, we can state:

Lemma 1. *Let $G = (\mathcal{I}, E)$ be the CDG of a feasible plan P . An ordering $\sigma'_{1:n}$ is incompatible with P iff there exists indices $u < v$ such that $(\sigma'_v, \sigma'_u) \in E$.*




Figure 3.4: A set of plans P_1, \dots, P_4 (top) and their dependency graphs G_1, \dots, G_4 (bottom). For any ordering beginning with A, there is at least one plan (P_1 or P_2) compatible with it. But for any ordering beginning with BDA, CB, CD, DAC, or DC, no plans are compatible.

In other words, a compatible ordering obeys all pairwise ordering constraints specified by the edges of the CDG.

3.4.4 Coverage verification

A key subroutine in our algorithm is to verify whether a set of packing plans P_1, \dots, P_m is compatible with all orderings in S_n , and if not, to generate a counterexample (i.e., incompatible ordering). An example is shown in Fig. 3.4. Let us reduce this to a combinatorial problem of validating whether a set of dependency graphs is compatible with all orderings, and call it DEPSET-COMPAT.

We present a recursive algorithm, which tries assigning each unassigned vertex v , and recurses on the subset of plans in which v is a root. There are two base cases:

1. There exists a vertex v that is not a root in any G_i . Then, any ordering that starts with v is a counterexample.
2. G_i has no edges for some plan P_i . The policy is feasible because P_i is compatible with all orderings.

To verify faster, we also perform a *singleton pruning* step: if a vertex v is a singleton (has no neighbors) in every G_1, \dots, G_m , then v can be safely ignored. This is because v can be assigned at any point without affecting dependencies.

The overall algorithm is given a vertex set \mathcal{I} and the edge sets E_1, \dots, E_m of the CDGs of P_1, \dots, P_m as input, and is listed in Alg. 5. The return value is either “all compatible” or a subsequence of \mathcal{I} that is incompatible with every dependency graph. Line 1 processes the first base case, and line 3 processes the second. Line 3 performs the singleton pruning step, and Lines 4–14 perform the recursion. Lines 5–10 compute the list E^v of dependency graphs that are compatible with assigning v at the current step, but with v is removed. In Line 13, the vertex v is prepended to the counterexample of a recursive call, because the counterexample is reached after assigning v .

A counterexample can often be found faster by ordering the vertices in Line 4 using a heuristic. Our approach sorts the vertices v by the number of plans compatible with the assignment of v (i.e., have v as a root).

Alg. 6 solves NDOP using this subroutine.

3.4.5 Quasi-online packing

Due to the need for shared transforms, QOP is not as amenable to elimination of orderings via compatibility verification. A naïve method for QOP would build a policy tree by enumerating all possible orders and ask for compatible plans.

Algorithm 5: Verify-CDG-Coverage($\mathcal{I}, (E_1, \dots, E_m)$)

input : a set of items \mathcal{I}

dependency graphs (E_1, \dots, E_m)

- 1 **if** *there exists* $v \in \mathcal{I}$ *that is not a root in any graph* (\mathcal{I}, E_i) **then return** “ v incompatible”;
- 2 **if** *any* E_i *is empty*, $i = 1, \dots, m$ **then return** “all compatible”;
- 3 Remove all vertices v from \mathcal{I} that are singletons in every graph (\mathcal{I}, E_i) ,
 $i = 1, \dots, m$;
- 4 **for** $v \in \mathcal{I}$ **do**
 - 5 $\mathcal{E}^v \leftarrow ()$;
 - 6 **for** $i = 1, \dots, m$ **do**
 - 7 **if** v *is a root in* E_i **then**
 - 8 Append \mathcal{E}_i to E^v , but with v removed;
 - 9 **end**
 - 10 $r \leftarrow$ Verify-CDG-Coverage($\mathcal{I}/\{v\}, \mathcal{E}^v$);
 - 11 **if** $r = “\sigma_{1:j}$ incompatible” **then**
 - 12 **return** “ $v, \sigma_{1:j}$ incompatible”
- 13 **end**
- 14 **return** “all compatible”

Algorithm 6: NDOP

input : a set of items \mathcal{I}

output: a solution set of plans \mathcal{P} , or “failure”

```
1  $\mathcal{E} \leftarrow$  empty-list;
2  $\mathcal{P} \leftarrow$  empty-list;
3 while true do
4    $r \leftarrow$  Verify-CDG-Coverage( $\mathcal{I}, \mathcal{E}$ );
5   if  $r =$  “all compatible” then return  $\mathcal{P}$ ;
6   Let  $\sigma_{1:k}$  be the incompatible ordering in  $r$ ;
7    $P \leftarrow$  Offline-Pack(nil, nil,  $\sigma_{1:k}$ );
8   if  $P =$  “failure” then return “failure”;
9   Add  $P$  to  $\mathcal{P}$ ;
10  Add  $CDG(P)$  to  $\mathcal{E}$ ;
11 end
```

Specifically, let N be a node in the policy tree at depth k , which is associated with the feasible plan $P = (\sigma_{1:n}, T_{1:n})$. For all non-placed items $\sigma'_{k+1} \notin \sigma_{1:k}$, we could call:

$$P' \leftarrow \text{Offline-Pack}(\sigma_{1:k}, P, \sigma'_{k+1}). \quad (3.7)$$

If P' = “failure”, then failure is returned. Otherwise, P' is associated with a new child of N in the tree corresponding to the choice σ'_{k+1} , and the search can proceed recursively. Note that if σ_{k+1} was already the $k + 1$ 'th item in P , replanning is unnecessary and we can just set $P' = P$. With this check, only $O(n!)$ calls to the offline planner are needed.

This procedure can be optimized by observing that all items that are roots of the dependency subgraph $CDG(\sigma_{k+1:n}, (T_{\sigma_{k+1}}, \dots, T_{\sigma_n}))$, can reuse P . In fact, all *combinations of roots* can reuse P . Moreover, once roots have been assigned, any newly created children can also reuse it.

To exploit this, our QOP planner performs a depth-first search while maintaining a list of plans \mathcal{P}_N compatible with $\sigma_{1:k}$ (i.e., all T_j match, for each fixed $j \in \sigma_{1:k}$). The search proceeds to enumerate children of $\sigma_{1:k}$. For each choice σ_{k+1} and child node C , if at least one plan in \mathcal{P} is compatible with σ_{k+1} , replanning is not performed, and the choice $T_{\sigma_{k+1}}$ is fixed. If multiple plans are compatible, the value of $T_{\sigma_{k+1}}$ that is compatible with the most plans is used. \mathcal{P}_C is then set to the set of plans in \mathcal{P}_N for which σ_{k+1} is a root of the dependency subgraph, and whose placement of σ_{k+1} matches $T_{\sigma_{k+1}}$. If no plan is compatible, then Offline-Pack is called as normal, and \mathcal{P}_C is set to contain only the newly generated plan P' . Moreover, we add P' to the sets \mathcal{P}_A for any ancestor of C . This enables subsequent siblings, siblings of parents, etc. to use P' and avoid additional planning. Pseudocode is given in Algs. 7 and 8.

3.4.6 Analysis

Here we show that NDOP inherits the completeness properties of the offline planner, but QOP is incomplete. Even when the offline planner is incomplete, when the NDOP or QOP result is not “failure”, the solution is correct. We also analyze the behavior of Verify-CDG-

Algorithm 7: QOP-Recurse(N)

input : policy tree node N

- 1 Let $\sigma_{1:k}$ be the sequence of packed items in N ;
- 2 Let \mathcal{P}_N be the set of compatible plans with N ;
- 3 Let P be any plan in \mathcal{P}_N , or *nil* if $\mathcal{P}_N = \emptyset$;
- 4 **if** all $\sigma_{k+1} \notin \sigma_{1:k}$ are roots of P **then return** “success”;
- 5 **for** all items $\sigma_{k+1} \notin \sigma_{1:k}$ **do**
 - 6 **if** no plan in \mathcal{P}_N is compatible with σ_{k+1} **then**
 - 7 $P' \leftarrow \text{Offline-Pack}(\sigma_{1:k}, P, \sigma_{k+1})$;
 - 8 **if** $P' = \text{“failure”}$ **then return** “failure”;
 - 9 $\mathcal{P}_C \leftarrow \{P'\}$;
 - 10 $C \leftarrow \text{add-child}(N, \sigma_{k+1}, \mathcal{P}_C)$;
 - 11 For all ancestors A of C , add P' to \mathcal{P}_A ;
 - 12 **else**
 - 13 Let $T_{\sigma_{k+1}}$ be the location compatible with the most plans in \mathcal{P}_N ;
 - 14 $\mathcal{P}_C \leftarrow \{P' \in P \mid P' \text{ is compatible with } T_{\sigma_{k+1}}\}$;
 - 15 $C \leftarrow \text{add-child}(N, \sigma_{k+1}, \mathcal{P}_C)$;
 - 16 **end**
 - 17 **if** QOP-*Recurse*(C, \mathcal{P}_C) fails **then return** “failure”;
- 18 **end**
- 19 **return** “success”

Algorithm 8: QOP()

- 20 $root \leftarrow \text{make-node}(\text{nil}, \emptyset)$;
- 21 **if** QOP-*Recurse*($root$) is successful **then return** $root$;
- 22 **else return** “failure”;




Figure 3.5: An example showing that QOP (Alg 8) is not necessarily complete even with a complete offline planner. (a) The first recursive call produces a feasible plan with A placed first. (b) Once item A is placed in the planned location, the plan is infeasible for order ACB, as shown in (c). On the other hand, if A was placed as in (d), a feasible QOP solution could result.

Coverage and demonstrate that it is NP-complete.

Correctness and completeness

NDOP inherits its completeness from the offline packing planner. To see this, first observe that whenever NDOP returns a solution (Line 6), this solution is correct, because for all orderings $\sigma_{1:n} \in S_n$, Alg. 5 has shown that the solution contains some plan that is compatible with $\sigma_{1:n}$. Now consider the case where NDOP returns “failure.” This can only occur when Offline-Pack returns failure for a partial ordering $\sigma_{1:k}$ (Line 7). If Offline-Pack is complete, then there is indeed no solution compatible with this ordering, and hence NDOP returns failure correctly. If it is incomplete, then NDOP may return failure incorrectly.

Assuming Verify-CDG-Coverage takes negligible time, the worst-case running time of NDOP occurs when all n items are stacked upon one another. In this case, all $n!$ possible orderings must be examined for feasibility.

Unlike NDOP, QOP is not necessarily complete even if the offline planner is complete. This is because the offline planner may commit early to a bad choice because assumes it has control over future item ordering, as illustrated in Fig. 3.5. However, if it does return a solution, then this solution is feasible even if a heuristic offline planner is used.

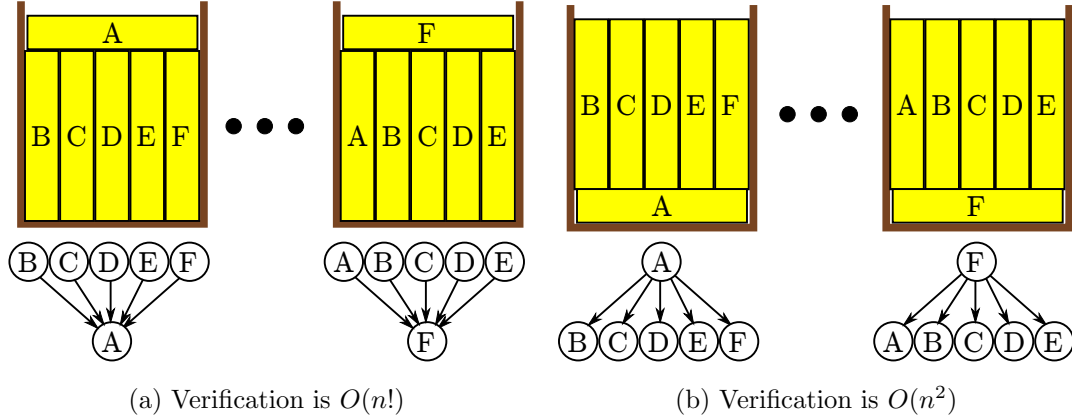


Figure 3.6: (a) Worst-case behavior of Verify-CDG-Coverage occurs in an instance with n plans, where $n - 1$ “books“ are stacked vertically with the n ’th book stacked horizontally on top. Every possible order of $k \leq n$ items is compatible with a set of $n - k + 1$ plans, and a recursion depth of n is required. (b) With a slightly different stacking, the dependency graphs are reversed. Only a depth 1 recursion is needed due to the singleton pruning step, so running time is polynomial.

DEPSET-COMPAT is NP-Complete

We observe that Verify-CDG-Coverage terminates extremely quickly in many cases, but can exhibit exponential behavior. An example is shown in Fig. 3.6.a, in which each of the $m = n$ plans has one item depending on all other items. At each level ℓ of the recursion tree, there are $n - \ell$ valid CDGs, and all $n - \ell$ vertices are valid. Hence, the function is called $O(n!)$ times. In fact, we prove the following theorem:

Theorem 2. *DEPSET-COMPAT is NP-complete.*

Proof. The proof is via polynomial time reduction from 3-SAT. A 3-SAT instance consists of n Boolean variables x_1, \dots, x_n and a logical expression in disjunctive normal form, consisting of m clauses

$$(y_{11} \vee y_{12} \vee y_{13}) \wedge \dots \wedge (y_{m1} \vee y_{m2} \vee y_{m3}) \tag{3.8}$$

where y_{ij} indicates either a variable or its negation, i.e., $y_{ij} = x_k$ or $y_{ij} = \neg x_k$. We transform

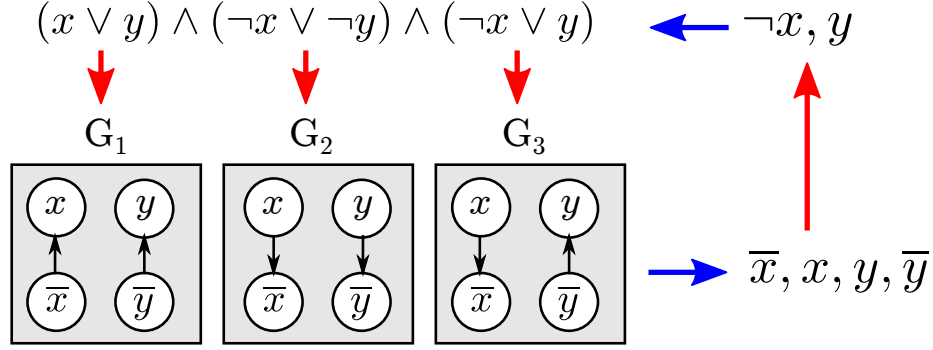


Figure 3.7: Illustrating the reduction from SAT. Each clause (upper left) is converted into a dependency graph (lower left), and a counterexample (lower right) ordering corresponds to a SAT solution (upper right).

any 3-SAT instance in this form into the complement of a DEPSET-COMPAT instance on $2n$ vertices and up to m dependency graphs. That is, when 3-SAT has a solution, the DEPSET-COMPAT version returns an incompatible ordering which corresponds to a 3-SAT solution, and when 3-SAT has no solution, the DEPSET-COMPAT version returns “all compatible.” This construction is illustrated in Fig. 3.7 for a 2-SAT instance.

Specifically, let $\mathcal{I} = \{v_1, \dots, v_n, \bar{v}_1, \dots, \bar{v}_n\}$ be the vertex set. Consider the i 'th conjunctive clause in (3.8). First, if the same variable and its negation appear in the same clause (e.g., $x_4 \vee \neg x_4 \vee x_6$), we drop the clause because it is satisfied via any assignment. Otherwise, we construct a dependency E_i as follows. If $y_{ij} = x_k$ for some k , construct an edge (\bar{v}_k, v_k) . If $y_{ij} = \neg x_k$, construct an edge (v_k, \bar{v}_k) . In the first case, this means that if this dependency is violated, then v_k will appear before \bar{v}_k in the ordering. In the second case, the reverse is true. This is repeated for each $j = 1, 2, 3$ and $i = 1, \dots, m$.

If the $\text{DEPSET-COMPAT}(\mathcal{I}, E_1, \dots, E_m)$ instance constructed in this way returns an incompatible ordering, we observe whether each v_k appears before \bar{v}_k . If so, we assign $x_k \leftarrow T$, and if not, $x_k \leftarrow F$. The variables x_1, \dots, x_n are then a solution to 3-SAT. This holds because in every clause $y_{i1} \vee y_{i2} \vee y_{i3}$, the dependency graph E_i is violated in such a way that makes the clause true.

Conversely, if the 3-SAT instance has a solution (x_1, \dots, x_n) , the DEPSET-COMPAT

instance has an incompatible ordering. It is constructed as follows: place v_k before $\overline{v_k}$ if $x_k = T$, and $\overline{v_k}$ before v_k if $x_k = F$. This ordering violates at least one constraint in each dependency graph.

Since each step in the reduction is polynomial time and NP-complete, DEPSET-COMPAT is NP-hard. It is also in NP, since a nondeterministic recursion could enumerate all possible orderings and check their validity in $O(mn)$ time. \square

What is interesting about this reduction is that DEPSET-COMPAT is hard even if restricted to seemingly easy classes of dependency graphs, e.g., separable, bipartite graphs with at most 3 dependencies! Experimentally, we have observed that DEPSET-COMPAT problems corresponding to hard 3-SAT problems (e.g., with clause-to-variable ratio of ~ 4.24 [Fre96]) also exhibit exponential complexity when solved via Alg. 5.

3.5 Planning Heuristics

Although DEPSET-COMPAT is NP-complete, computation time of NDOP and QOP is dominated by time spent in the offline planner, because each plan requires searching over 6D object pose. The number of plans requested, and hence overall running time, is greatly dependent on the number of orderings compatible with previous offline plans. Hence, it would be beneficial if the offline planner would generate packing plans that maximize compatibility. We employ some heuristics that speed up the approach in common scenarios.

3.5.1 Dependency minimization heuristic

Constructive packing chooses an item’s location based on certain placement heuristics, such as deepest-bottom-left-first (DBLF) [WGC⁺10], or heightmap minimization (HM) [WH19b, WGC⁺10], to maximize packing density. For nondeterministic packing, we would like to generate plans with few dependencies. We introduce a *dependency count* (DC) heuristic that measures the number of items underneath the item at the given placement (i.e., number




Figure 3.8: The dimensions (in inches) of containers 1–5 used in our experiments, in order of increasing length + girth.

of ancestor nodes in the CDG). Our implementation uses a heuristic that is a weighted sum of HM and DC.

3.5.2 Matching prior placements

In QOP it is beneficial for the offline planner to place as many “free” items (i.e., those not in $\sigma_{1:j}^{fixed}$ or $\sigma_{j+1:k}^{next}$) as possible in the same location as the prior plan, since this will maximize the likelihood that the plan is compatible with other branches in the search tree. To implement this heuristic, when packing a free item, the location in P^{prior} is checked for feasibility before any other locations are tested.

3.5.3 Container optimization heuristics

During container optimization, it is helpful to limit exponential growth in running time by replacing the infinite loop in Line 3 of Alg. 6 with a fixed number of iterations, or break the recursion of QOP after the policy graph has grown too large. As the containers grow wider / longer, the number of dependencies decreases because all items will be packable in fewer layers. With a large enough container, all items are packable in a single layer. Hence, if there exists a sufficiently large container, the optimization version will always terminate with a feasible, but possibly suboptimal solution.

Table 3.1: NDOP container optimization results

Items	Success (%)	Time (mean / max, s)	# planner calls (mean / max)
2–5	99.3	73.0 / 1,813	1.3 / 9
5	96.9	94.4 / 1,417	1.6 / 14
10	64.0	1,048 / 33,300	5.2 / 118

3.6 Experiments

Our experiments test the NDOP and QOP algorithms with random item sets of 3D scanned objects from the APC 2015 [Rut15] and YCB [CSB⁺17] datasets (94 objects total). The containers used in these experiments are the five boxes used in the Amazon Robotics Challenge 2017 (Fig. 3.8), and are sorted by the length + girth metric (length + 2×width + 2×height), a commonly-used shipping measurement. All experiments were performed on an Amazon EC2 m5d.12xlarge instance.

Our experiments use three testing datasets in which the item sets have different size: a) typical shopping carts of 2–5 items, b) large sets of 5 items, and c) stress tests with 10 items. For each category, we generate 1,000 random item sets by drawing items at random, and verifying with an offline planner that there exists a feasible packing in one of the five containers. In the 2–5 category, 250 item sets of each size are included. The results for the NDOP planner, running in container optimization mode, are summarized in Tab. 3.1, and Fig. 3.9 illustrates a solution. As might be expected, the running time and the number of offline planner calls increases with the number of items, but we do not observe the exponential running time of worst-case instances. Other experiments suggest the dependency minimization heuristic reduces mean and maximum running times by approximately 20% and 50%, respectively.

Observe also that the success rate drops with increasing numbers of items, as the 10-item offline plans tend to be tightly packed even in the largest container. Note that it is

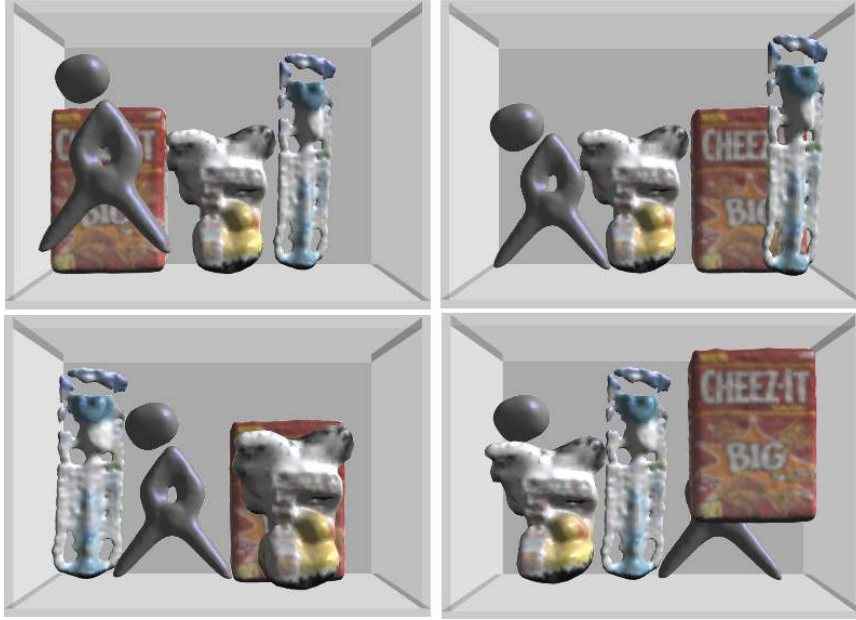


Figure 3.9: An NDOP solution for packing 5 items in container 5.

not known whether an NDOP solution exists in these instances, so we cannot determine whether the solver is failing incorrectly. Fig. 3.10 shows the distribution of the minimum-cost container found. In cases with 5 or fewer items, NDOP often successfully packs in the same box as the offline planner. Observe that with 5 items, approximately 5% of test cases require container 5, even though they can be packed offline in containers 1–4. Fig. 3.11 illustrates an example of such a case.

Performance results for QOP in container 5 are given in Tab. 3.2, with a representative solution shown in Fig. 3.12. Up to 5 items, the success rates are quite similar to NDOP, but the maximum running times and number of offline planner calls tend to be significantly larger. Other experiments suggest that employing the matching heuristic improves average and maximum running time by over 50%, which explains the surprising result that QOP is faster than NDOP on average. QOP struggles with 10 items, with a long-tailed distribution: 24 instances could not be solved within a 24-hour cutoff.




Figure 3.10: Distribution of the solution container for offline / NDOP container optimization and various itemset sizes. Most small instances can be packed in any order, but with more items the variability in order requires larger boxes.




Figure 3.11: A failure case for packing 5 items in container 3. The offline planner solves for the arrival order in (a) but fails on the order in (b) because there is insufficient remaining space for the fifth item.

Table 3.2: QOP planning results in container 5

Items	Success (%)	Time (mean / max, s)	# planner calls (mean / max)
2-5	99.4	22.4 / 1,520	1.4 / 43
5	97.0	65.1 / 5,800	2.1 / 46
10	43.7	2,850 / 85,478	45.8 / 5,363




Figure 3.12: A QOP solution for 4 items in container 5. Due to the matching heuristic, only four plans are needed (one for each item placed last).

Chapter 4

In-hand Object Scanning via RGB-D Video Segmentation

The chapter introduces useful tools to gather 3D shape information for an arbitrary object, we proposes a technique for building full 3D models of an arbitrary object via in-hand manipulations, in which the object is reoriented in front of a video camera with multiple grasps and regrasps. 3D models with colors and accurate geometry are known to increase robustness in a robot’s perception and manipulation. Therefore, building accurate 3D models of task-relevant objects and recording their identities could greatly benefit the autonomy of many robotic tasks in the environment.

In-hand object scanning is a convenient and inexpensive way to model a hand-sized object, yet many significant challenges exist in tracking the manipulated object due to fast movement, rapid appearance changes, and occlusions. We proposes a novel video-segmentation-based object tracking algorithm that tracks arbitrary in-hand objects more effectively than existing techniques. It also describes a novel RGB-D in-hand object manipulation dataset consisting of several common household objects. Experiments show that the new method achieves 6% increase in accuracy compared to top performing video tracking algorithms and results in noticeably higher quality reconstructed models. Moreover, testing with a novice user on a set of 200 objects demonstrates relatively rapid construction of complete 3D object models.¹

¹This chapter is reproduced from Fan Wang and Kris Hauser, “In-hand Object Scanning via RGB-D Video Segmentation,” in 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 2019, pp. 3296-3302.

4.1 Introduction

The ability to build 3D models of the geometry and appearance of real-world objects has long been essential in many application scenarios such as art and design, manufacturing and packaging, augmented reality, and robotics. While early 3D model acquisition relied on multiple calibrated scanners or accurate motion apparatuses [RHHL02, BR02, KHRF11], the last decade has seen the introduction of low-cost RGB-D cameras like Microsoft Kinect and advances in feature detection and automatic feature matching algorithms that have made 3D object model acquisition easier than ever.

Among various 3D scanning techniques, in-hand scanning has attracted much research attention [TG17, PASV15, WWLG09, PKA15]. Methods that require an object to be stationary during the scan inherently carry the missing side problem [BR02], while in-hand scanning can reveal all sides of the object, making it possible to reconstruct the complete object without multiple model alignment in post-processing. The human hand is capable of intricate motion when manipulating objects through gripping, grasping and turning, so this is a natural way of presenting an object for scanning. It is also low-cost and flexible since it does not require expensive equipment and calibration. Many reconstruction pipelines from object-hand interaction require only an off-the-shelf RGB-D camera [TG17, PASV15].

Despite these attractive characteristics, 3D object modeling from in-hand object manipulation remains a challenging problem. To obtain a high-quality object model, the target object should be accurately separated from non-target objects. This separation is difficult since hands may be similar to the target object in terms of color, motion, or appearance, and furthermore, the target object itself undergoes significant changes in a short period of time.

This chapter presents a 3D model acquisition pipeline from in-hand interaction. The pipeline includes a novel object tracking technique and a set of reconstruction and post-processing procedures. With this pipeline, a non-expert can scan arbitrary objects with only a hand-held single RGB-D camera and light manual annotation.

To evaluate our work, we constructed a public in-hand object manipulation dataset



Figure 4.1: Some successfully reconstructed objects from a 200-item test set, scanned by a novice user of our system.

consisting of 13 objects from the publicly available YCB object set [CWS⁺15] being manipulated by hand in front of an RGB-D camera. Experiments on this dataset demonstrate that our method outperforms many state-of-the-art video segmentation algorithms in terms of tracking performance and results in higher quality 3D reconstructed models. We also asked a novice user to use our pipeline to scan 200 arbitrary items, requiring approximately 2.5 minutes of manual effort per object including scanning and annotation, and demonstrates the capability to produce high-quality reconstructed models, shown in Fig. 4.1.

4.2 Related Work

An object scanning pipeline commonly consists of object-background segmentation, multi-view registration, and registration refinements steps. Several in-hand scanning systems have been presented [WLG08, WWLG11, TG17, RHHL02, PKA15]. Weise et al.[WLG08] demonstrated real-time system reconstruction of textured objects from in-hand manipulation using coarse to fine registration and online loop closure. Tzionas et al.[TG17] showed the possibility of reconstructing featureless objects from hand motion cues, by minimizing objective functions that consist of both visual correspondences and contact correspondences.

While the registration and refinement aspects of an object scanning system have been extensively studied[WLG08, WWLG09, RHHL02, BR02, TG17], little work has been done to address the unique challenges present in foreground - background segmentation in an in-hand manipulation task, despite its importance being commonly acknowledged [WWLG09, KHRF11, SMZ⁺16, PKA15].

Existing in-hand scanning systems address the segmentation task with simple approaches. Hand-object segmentation are commonly addressed by skin pixel identification based on statistically color model [WLG08, TG17], or by having operators wear black gloves [WWLG09, WWLG11, RHHL02]. To remove other unwanted backgrounds, depth thresholding and black backgrounds are commonly used [TG17, RHHL02]. A more sophisticated method is to perform segmentation in phase image as in [WLG08] to remove background

that remains stationary.

However, those methods are not robust enough to achieve highly accurate object-background segmentation. Color based skin-pixel segmentation, for example, can result in a high false negative rate with hands under different lighting conditions and a high false positive rate when objects being manipulated are close in color to human hands, as illustrated in Fig. 4.2. The phase image segmentation cannot remove non-stationary background objects such as hands and arms. The black glove method might be the most robust among all, but nonetheless requires only the hand and the object to be visible during the entire scanning, which can be difficult for the operator to accomplish.

Our pipeline uses similar registration approaches to prior work, but to address hand-object separation with more generality and to higher accuracy, we solve a semi-supervised binary video segmentation problem, tailored for the hand-object segmentation task. Our method can separate objects accurately from all unwanted backgrounds, including hands of various skin tones, as well as other stationary or non-stationary objects not of interest.

Recent research on video segmentation-based object tracking can be roughly categorized into unsupervised, semi-supervised, and human-in-the-loop approaches.

Unsupervised methods are fully automatic and require no human effort during run time [LTHF17, XL16, WSP15, CTWY17]. While this approach is preferred when real-time tracking is necessary, a challenge is to propose a foreground hypothesis without specific knowledge on the actual object to track. Proposals are commonly formed from pre-learned models of “object-like” regions using static cues such as convex regions, pixel group similarities, or classification based on learning [HPL17, RT16, KMV⁺16] as well as motion cues [TYB16, CTWY17], given that rigid foreground objects should move with coherent motion different from the background. Unsupervised segmentation for object scanning is difficult because pre-training is unlikely to apply to the large variety of arbitrary objects to be scanned, and motion cues from human hand movement can mislead the algorithm.

Semi-supervised and human-in-the-loop methods require varying amount of human interaction to specify the object of interest for higher accuracy [DAV]. In semi-supervised



Figure 4.2: Skin regions identified by skin-pixel based segmentation described in [VSA03]. This method cannot detect shaded areas of hand as well as accessories on hand (jewelry, nail polish, etc.). It also falsely includes object regions with colors close to human hands.

methods, the human provides an initial ground truth, typically for the first frame of a video sequence. This ground truth could be a segment mask [CMPT⁺17, CTWY17] or user scribbles annotating foreground and background [RKB04, BK04]. The algorithm then propagates this annotation segment to the unlabeled frames. Some semi-supervised methods use the first frame for fine-tuning a learned object model [CMPT⁺17, PKB⁺17, VL17] and then perform per-frame segmentation. However, most semi-supervised methods use a propagation scheme to propagate the initial labels to successive frames [TYB16, PS14, CTWY17].

4.3 In-hand Object Reconstruction Pipeline

Our pipeline takes an RGB-D video as input, in which a user rotates the object in front of the camera. The user can perform multiple regrasps to reveal all sides of the object. Our method then uses a novel semi-supervised video segmentation technique, called BackFlow, to segment the object from the hands and background. These object segments are then registered to a global reference frame, and the overall model is post-processed to reduce noise.




Figure 4.3: The new sides problem poses a challenge for video segmentation algorithms that rely on coherent object appearance. Tracking a can of tomato soup with a hierarchical graph-based method [GKHE10] fails to grow the segmentation to newly appeared top of the can.

4.3.1 BackFlow Video Segmentation Method

BackFlow is designed to handle many unique challenges for hand-object segmentation in RGB-D videos. For example, the **new sides problem** occurs because objects being manipulated quickly undergo significant shape and appearance changes. These changes are difficult to predict without prior knowledge of the object. Conventional video segmentation methods have difficulty maintaining tracking (Fig. 4.3). **Object-hand appearance similarity** is also quite common. Although the appearance of human hands varies from person to person, various skin-tone colors are commonly represented on the packaging. Additionally, the hand could also be wearing nail polishes or jewelry, which further complicates the identification of hand from the object. **Occlusion and large deformation** are especially prominent with human hands. Hands can occlude large portions of the object, and deform quickly and significantly. **Finger-object motion similarity** also tends to confuse algorithms that rely on distinct motion cues.

BackFlow derives new object appearance hypotheses by deduction, rather than through foreground coherence. It preferentially tracks background pixels and derives foreground accordingly through a graph-based framework. Essentially, it assumes that what is not background is likely foreground, and it attempts to track and detect the disappearance and emergence of the background. The rationale behind this approach is that hands and

background objects exhibit stronger spatiotemporal coherence compared with the target object, and therefore are the better candidate to be tracked.

BackFlow requires no prior training and forms initial object proposals based on color GMMs learned in the first segment. The proposals are adjusted during run-time with more available frames. This method is most effective when the appearance and shape of the foreground object evolve quickly while the background remains relatively stable.

The four main steps of BackFlow are described below:

Preprocessing and initialization

BackFlow first eliminates unnecessary background tracking. A depth cut-off on the color images is performed that eliminates areas more than 1 m away from the camera. Then it requires a human to initialize *sure background* pixels, either by providing background scribbles or a segment mask. These annotations serve as hard constraints in graph-cutting for the first frame.

Background pixel propagation

Next, we use optical flow [Far03] to propagate raw background pixels between two successive frames. The flow is cross-examined by performing backward flow, and background pixels with the inconsistent flow are dropped. Since most optical flow inconsistencies occur near the motion edge, this helps prevent background labels from drifting into the foreground.

Superpixels are employed to guide the background propagation and to increase robustness against drift. Specifically, we perform SLIC segmentation [ASS⁺12] with $N = 20,000$ superpixels on frames with 640×480 resolution. When propagating background pixels from frame t to frame $t+1$, all frame t superpixels that contain raw background pixels are marked as the background superpixels. We track the flow of all pixels within the background superpixels to the second frame and label frame $t+1$ superpixels that receive enough propagated background pixels as new background superpixels. The use of superpixels provides additional color information and prevents drifts into neighboring groupings that differ in color.

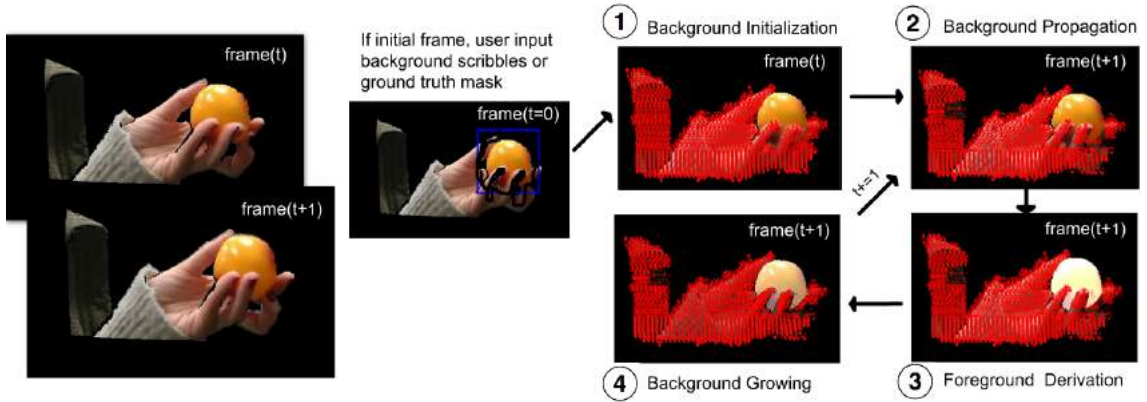


Figure 4.4: BackFlow takes user annotation of background or ground truth segment on the first frame and propagates labels as follows. 1) Initialize background labels (superpixels) from user input. 2) Propagate background labels from frame t to frame $t+1$ through optical flow. 3) Perform graph-based foreground derivation. 4) Remove background labels in segmented foreground and grow background labels to neighbourhood superpixels.

This measure also replenishes background labels within a superpixel under the assumption that all pixels within a superpixel likely belong to the same label.

Segmentation with relaxed foreground derivation

Step 3 performs graph-based segmentation after the background pixels have been propagated to the current frame. Pixels form vertices connected to neighbors and both background and foreground labels by weighted edges. The segmentation is formulated as an energy minimization problem in which the energy function has the form:

$$E(L) = \sum_{p \in \mathcal{P}} D_p(L_p) + \gamma \sum_{(p,q) \in \mathcal{N}} V_{(p,q)}(L_p, L_q), \quad (4.1)$$

where $L = \{L_p | p \in \mathcal{P}\}$ is a labeling of image \mathcal{P} , $D_p(\cdot)$ is the data penalty function assigning labels to the vertex, $V_{(p,q)}$ is the smoothness term, and \mathcal{N} is a set of all pairs of neighboring pixels [BK04]. γ is a parameter chosen to balance data and smoothness terms. The

minimum-energy cut of the graph is therefore a partition of the vertices into two disjoint sets, solved by min-cut/max-flow algorithms [BK04].

We use Gaussian Mixture Models (GMMs) to describe background and foreground distributions similar to [RKB04]. We modify standard graph segmentation algorithms to encourage foreground propagation and to rely more heavily on motion cues in the presence of drastic appearance changes.

Background GMMs with $K=8$ components are learned on the first frame and used throughout the video sequence, and foreground GMMs are learned on a sample of raw pixels in the previous foreground segments. Specifically, we store past foreground pixels, and after each new frame is segmented, we downsample the foreground set by 90% and add 10% new foreground pixels; this is to ensure a stable background proposal and a smoothly transitioning foreground proposal as the object evolves.

In BackFlow, we formulate initial data penalty terms similar to Grabcut [RKB04], which indicate the log likelihood of a pixel belonging to the foreground and background GMMs, respectively:

$$\begin{aligned} \tilde{D}(\alpha_p, k_p, \underline{\theta}, z_p) = & -\log \pi(\alpha_p, k_p) + \frac{1}{2} \log \det \sum (\alpha_p, k_p) \\ & + \frac{1}{2} [z_p - \mu(\alpha_p, k_p)]^\top \sum (\alpha_p, k_p)^{-1} [z_p - \mu(\alpha_p, k_p)], \end{aligned} \quad (4.2)$$

where α_p indicates background or foreground assignments and has value of 0 or 1 in hard segmentation. $\underline{\theta}$ stands for the GMM parameters. k_p is the GMM component assigned to the pixel and $k_p \in \{1, \dots, K\}$ where K is the number of components in the GMMs. z_p the a pixel color in RGB colour space. The weights π , mean μ and covariance Σ are the parameters for the Gaussians from the foreground and background distributions.

The data cost for each node is by default $D \leftarrow \tilde{D}$, except for the following exceptions. For nodes already associated with background labels, we assign D to a high constant penalty to be connected to the foreground and zero penalty to be connected to the background. For unlabeled nodes, we further add a comparison step in which we equalize data penalties assigned to both labels if a pixel only has a slightly lower probability belonging to the foreground than the background proposal. The background-foreground assignment of such

node is therefore determined by the assignments of neighborhood pixels that have a high interacting potential to the node. Specifically, if

$$\tilde{D}(\alpha_{p=1}, k_p, \underline{\theta}, z_p) > \tilde{D}(\alpha_{p=0}, k_p, \underline{\theta}, z_p), \quad (4.3)$$

and

$$\tilde{D}(\alpha_{p=1}, k_p, \underline{\theta}, z_p) - \tilde{D}(\alpha_{p=0}, k_p, \underline{\theta}, z_p) < N, \quad (4.4)$$

then

$$D(\alpha_{p=1}, k_p, \underline{\theta}, z_p) \leftarrow \tilde{D}(\alpha_{p=0}, k_p, \underline{\theta}, z_p), \quad (4.5)$$

where N is the relaxation constant. Experiments find that the segmentation accuracy is not greatly impacted by the value of N for $8 < N < 16$.

The introduction of the relaxation constant encourages the inclusion of the newly appeared foreground that is adjacent to the existing foreground. While this measure would otherwise introduce false positives if the background is not densely labeled, the majority of the background area is already correctly labeled. We also observed that the difference in data penalties for newly appeared background (e.g., hands, arms) are usually much greater than the relaxation constant.

A commonly used smoothness term (e.g., in Grabcut) is:

$$V(\underline{\alpha}, z) = \sum_{(m,n) \in \mathbf{C}} I[\alpha_n \neq \alpha_m] \exp(-\beta \|Z_m - Z_n\|^2)$$

z is the given image data, $\underline{\alpha}$ is the background or foreground assignments for each pixel, \mathbf{C} is the set of all neighbouring pixels, $\|Z_m - Z_n\|^2$ is the Euclidean distance between neighbouring pixels in RGB color space. β measures the inverse expectation of contrast within the image.

In BackFlow we introduce motion cues into the smoothness term. To encourage consistent labeling for areas with coherent motions, we use optical flow along with color dissimilarities for smoothness term calculations. We convert the flow vectors into images in HSV color space where the Hue and Value are computed from the angle and magnitude of the




Figure 4.5: Background labels (marked as red dots) grow gradually into the previously unseen palm while maintaining separation from the foreground object.

flow. The smoothness penalties are then taken as the smaller of the discontinuity in color and flow:

$$V(\underline{\alpha}, z) = \sum_{(m,n) \in \mathbf{C}} I[\alpha_n \neq \alpha_m] \min(\exp(-\beta_z \|Z_m - z_n\|^2), \exp(-\beta_f \|F_m - F_n\|^2)).$$

The expectation value β_z and β_f are calculated respectively for RGB image and the flow image. In our implementation we use $\gamma = 50$ and $N = 10$. The cut is then performed using Min-Cut/Max-Flow algorithms. The largest connected component in the result becomes the foreground segment.

Removing and growing background labels

Finally, we extend background tracking into the newly appeared background. After segmentation of each frame, all superpixels that are spatially adjacent to the existing background and are not segmented as foreground become background superpixels and will be tracked in the next frame as shown in Fig. 4.5.

4.3.2 3D Reconstruction

The segmented frames are registered to produce the final model. To avoid redundant information and bad frames with severe motion blur/noise, keyframes are selected for registration. Pairwise registration between keyframes is performed using sparse feature-point

matching followed by dense Iterative Closest Point (ICP) [BM92] alignment. Finally, we run pose-graph alignment and postprocessing.

By default, we select every 10th frame as the keyframe if it either 1) matches enough feature points from the previous keyframe or 2) its point cloud can be matched above a certain percentage with the previous keyframe with ICP. Otherwise, we try to match the previous keyframe with the 9th frame, 8th frame ... until a match is found. If no match exists, we will select the 10th frame and continue with the process.

For pairwise registration we extract SIFT (Scale-Invariant Feature Transform) [Low04a] features from consecutive keyframes and match feature pairs through a 2D RANSAC (Random sample consensus) homography process, and use the depth correspondence of the matched feature point to find the best 3D rigid body transformation between the two frames. If insufficient matching features are found, dense ICP is performed between two point clouds to estimate the transformation. We use the colored point cloud registration by Park et al. [PZK17] implemented in Open3D [ZPK18]. This ICP variant jointly optimizes both point-to-plane distances and minimizes the error of projected color. Pose graph optimization is used to obtain the final registration transforms in the global frame, with methods described in [KGS⁺11].

For postprocessing, the aligned colored point clouds are stored in voxel grids of dimension 2mm, and each voxel is optimized to reject color outliers, and grids with less than 2 points are not used. Registered point clouds then go through Poisson surface reconstruction [KH13] to create watertight surfaces, and isolated components whose diameter is smaller than 10% of the mesh diameter are removed to achieve the final models.

4.4 Experiments

4.4.1 RGB-D in-hand object manipulation dataset

Although there are several popular video segmentation datasets, there is no existing in-hand object scanning RGB-D video dataset with dense pixel level annotations. To systematically




Figure 4.6: RGB-D in-hand object manipulation dataset contains RGB-D video of 13 non-duplicate items from YCB food category, including Pringles, mustard, Cheez-It, sugar, Spam, tomato soup, banana, pear, plum, strawberry, orange, lemon and Jell-O.

evaluate our methods and to aid the research community, we introduce the RGB-D In-hand Object Manipulation Dataset (Fig. 4.6). This dataset contains 13 sequences of in-hand manipulation of objects from the YCB object set [CWS⁺15], recorded with an Intel RealSense SR300 RGB-D camera. Each sequence ranges from 300 to 800 frames in length (filmed at 30 fps) and contains in-hand manipulation of the objects revealing all sides. The dataset is complete with color images, depth cut-off images, and depth mages. The pixel-wise annotation aligned to the depth cut off image is provided every 10 frames. The dataset and results can be downloaded at the RGB-D In-hand Object Manipulation website at <https://www.rgbdinhandmanipulation.com>.

The performance of BackFlow is compared to other state-of-the-art segmentation methods in Tab. 5.2. OSVOS is a semi-supervised fully-convolutional neural network (FCNN) that was one of the top-performers on DAVIS 2017 [CMPT⁺17]. SFL indicates the un-

Table 4.1: Segmentation Results

	OSVOS	SFL	HVS	SPBS	BackFlow
IoU	87.97%	29.61%	52.05%	31.52%	93.52%
FP	4.17%	50.59%	1.7%	49.99%	2.12%
FN	7.85%	13.90%	46.29%	18.48%	4.36%

**Figure 4.7:** Per sequence accuracy (IoU) comparison of BackFlow to other state-of-the-art video segmentation methods.

supervised FCNN technique SegFlow [CTWY17]. Both OSVOS and SFL were pre-trained on the RGB-D Object Dataset [LBRF11] while OSVOS was additionally fine-tuned on the initial segment of each sequence. HVS is the hierarchical graph-based video segmentation of [GKHE10]. SPBS indicates skin-pixel based segmentation, which is a GMM method described in [VSA03], trained on the UCI skin segmentation dataset [Bha].

These experiments separate each object handling sequence into sub-sequences of 100 images. A ground truth annotation is given on the first frame of each sub-sequence, and the segmentation accuracy is tested on the remaining 99 frames. We measure the intersection over union (IoU), areas of false positives over union (FP) and areas of false negatives over union (FN). BackFlow outperforms other state-of-the-art approaches by almost 6% percent in overall accuracy and achieves lower false positive and false negative rate. Per-sequence results are shown in Fig. 4.7.

The reconstructed models of several objects from the dataset are shown in Figure 4.8.



Figure 4.8: Reconstructed models from BackFlow segmentation results.

The textured (color or geometry) objects are reconstructed complete with all sides, except when the surface material (e.g., black, highly reflective) cannot be captured by a depth camera. Our reconstruction pipeline currently cannot obtain a complete model with symmetrical and texture-less objects, such as pear and orange, and the registered point clouds appear as incomplete "shell".

Fig. 4.9 shows some example of unprocessed point cloud reconstruction from BackFlow against the next-best performer OSVOS. To qualitatively measure the quality of the reconstructed models, we further compare all 13 registered point clouds constructed from BackFlow and OSVOS to the ground truth Poisson meshes provided in the YCB dataset. The evaluation metric used is the one-sided Hausdorff Distance implemented in MeshLab. To measure it, the registered point cloud and the YCB mesh are aligned using manual key



Figure 4.9: Comparison of reconstructions from OSVOS and BackFlow.

points matching (we use the same rigid transformation for both registered point clouds), then the point cloud is sampled to 20,000 points and for each sample, the Hausdorff Distance to the closest point over the YCB mesh is computed.

Shown in Tab. 4.2, the Hausdorff distance measurements for BackFlow reconstructed models are around 50% smaller than OSVOS in all categories including max distance, mean distance, and RMS. The max distance in BackFlow is caused by failure to segment out a human finger in parts of the banana sequence, the artifact is reflected on the reconstruct

Table 4.2: Hausdorff Distance measurements w.r.t. bounding box size of the reconstructed models

	Max distance	Mean distance	RMS
BackFlow	0.195	0.014	0.017
OSVOS	0.427	0.024	0.031

(see Fig. 4.8), yet the max distance is still much smaller than that of OSVOS. While BackFlow only improves segmentation performance by 6%, the resulting reconstructions are of significantly better quality, as these segments maintain a better consistency and have a lower false positive rate.

4.4.2 Novice scanning of many items

To evaluate the ease of use and robustness of the pipeline when handled by a novice user, we presented our software to a volunteer who is unfamiliar with 3D object modeling. 200 grocery items were purchased satisfying the following criteria: 1) the size of the object should be easily manipulated with a single hand, 2) the object does not contain large black or transparent areas, which cannot be captured with a structured light RGB-D camera, 3) the object is textured or non-symmetric, 4) objects should not significantly deform during manipulation.

The volunteer was taught how to use the software and was told to reveal all sides of the object. No other instruction on how to manipulate the objects was given. We record the total amount of time needed, including the instruction time, and several small breaks, for the volunteer to finish scanning the objects. The volunteer then annotates every 100 frames of the scan sequence as initial labels.

On average, the volunteer took 102 seconds to scan each item, plus another 36 seconds for labeling (averaging 4.9 frames per item). Segmentation and reconstructions then automatically run to generate object reconstructions. The total labor time used for recon-

structuring 200 models is 7 hours and 40 minutes, or 1 fps for segmentation and reconstruction on a CPU computer.

As shown in Fig. 4.1, the pipeline is robust enough to generate many visually high-quality reconstructions with no hands or any other background objects found in any reconstructed model. There are also failed reconstructions, which are most frequently caused by misaligned frames during registration. Better registration techniques are likely to improve the results further.

Chapter 5

Systems and Analysis for Robust Robotic Packing

While the offline packing algorithm are guaranteed to be feasible under perfect modeling and precise execution, but making packing reliable enough on a actual physical system remains a major challenge due to numerous sources of system uncertainties. This chapter studies how errors in vision, grasping and modeling give cause cascading errors in the overall robotic packing pipeline. A systematic evaluation on a physical packing testbed studies the sources of error and models their magnitude. Experiments are conducted to quantify the impact of such errors on the overall packing success rate in Monte Carlo simulation and on the physical testbed. It also presents strategies for improving the robustness of robotic packing: 1) using conservative planning to ensure feasibility under uncertain ranges of model parameters, and 2) employ closed-loop vision and manipulation so that a robot can dynamically react to errors and correct for them. Empirical results demonstrate that the closed-loop packing and robust planning increase the overall packing success rate by 15%, from 83% to 98%, when compared to the baseline system. ¹

¹This chapter is reproduced from Fan Wang and Kris Hauser, “Systems and Analysis for Robust Robotic Packing,” in submission to IEEE Transactions on Robotics (T-RO). This work is supported by an Amazon Research Award.

5.1 Introduction

Small disturbances can cause serious, cascading failures, particularly if the packing arrangement is dense: if an object unexpectedly shifts during execution, it will likely leave insufficient space for remaining items in the selected container. These failures could cause damages to the items, since an industrial robot is stiffly geared and blind execution of pre-planned paths can exert very high forces when an unexpected obstacle is encountered. Such failures are also difficult and time-consuming to recover from, e.g., remove all placed items in the old box and replace them in a larger box.

In this chapter we analyse various sources of uncertainty inherent to automated packing systems, and propose two strategies for improving the reliability of packing under such uncertainties. To the authors' knowledge, this is the first complete robot packing system that is able to handle irregular objects.

The system uses state-of-the-art perception and planning components. The system picks from a set of arbitrarily shaped objects on a flat picking area and packs them in a shipping box. We use a planning algorithm that produces robot packing plans [WH19b], assuming a given location and dimensions of the box, 3D models of each object to pack, and initial object locations. We produce this input using a calibration setup step and a vision system to perform object detection and pose estimation. The robot then executes grasping and stowing motion to place each object into the container as planned.

The uncertainties in the system include object modeling errors, such as incorrectly estimated object shapes, poses, deformation, etc., manipulation errors, such as failed grasps, object reorientation during grasping, as well as other system-level uncertainties such as calibration error. We develop two strategies to mitigate the impact of uncertainties and increase the robustness of the packing execution. The first strategy is to plan conservatively during the offline planning phase to ensure execution feasibility under positioning error. The second strategy is to employ closed-loop vision and manipulation so that a robot can dynamically react to errors and correct them. Specifically, the system re-senses the object pile after each placement using visual and depth information and then replans the best

subsequent placements if the pile deviates from the plan significantly.

We perform exhaustive testing in both Monte Carlo simulation as well as on a physical robot, under different packing strategies. In Monte Carlo simulation, we test the packing success rate while systematically introduce positioning errors of arbitrary magnitudes on thousands of object sets. We have also tested those strategies on the physical platform (Fig. 5.1). Empirically, the execution success rate on the physical platform with an open-loop baseline system is quite low, at 83%, due to multiple sources of uncertainty causing cascading errors. The proposed strategies for robust robot packing raise the packing success rate to 98%.

5.2 Related work

Recent years have seen research that is more comprehensive in addressing practical constraints encountered in warehouse packing automation. For example, Ye et al. [YR18] develop a real-time state estimation system that can recover the pose and contact formation of a rectilinear object in grasp relative to its environment. By fusing visual and force sensing, they demonstrate the application of inserting an object picked by a suction cup into a tight space. This work, however, does not focus on planning, but rather on the fitting of a single item into a confined space.

The work arguably closest to ours is by Shome et al. [STS⁺19]. They designed a complete robot packing system using minimalistic hardware stack of a single robot with suction gripper equipped with RGB-D cameras, similar to ours. They have also employed different techniques to achieve robust packing that minimizes failure conditions, One example of such technique is using the suction gripper as a push finger that executes three pre-defined key manipulation primitives to tight packing in against perception and positioning errors. The key difference between their system and ours is they test on uniformly shaped cuboid objects, where the planning aspect is relatively simple. Ours focuses on tight and stable placement plannings for arbitrarily shaped objects that can be convex or concave, which is

a much harder planning problem.

We use the offline planner from our previous work that produces robot-packable object placements. Compare with the other methods mentioned, this algorithm accepts arbitrary shaped objects and provides verification that a container fits all the objects, therefore more suitable for a warehouse packing scenarios where a set of customer ordered items need to be packaged into a single shipping box. The other rationale behind using the method is that compared to compliant motion-based packing as in [STS⁺19, YR18], which intentionally introduces collision to the object during manipulation to force tight placement, planning for high volume density packing likely results in less damage to the object during packing execution.

5.3 System setup and error analysis

In this section, we first introduce our hardware and software setup is fairly representative of current state-of-the-art sensing, perception, and planning for warehouse manipulation. Then we identify the sources of error in such a setup that may lead to failed packing, and propose two strategies for mitigating these errors.

5.3.1 Experimental setup

A 6R industrial robot (Universal Robot UR5e) is equipped with a suction gripper (Robotiq EPick) and a compliant suction cup (Fig. 5.1). One RGB-D camera (Realsense SR300) and two high-resolution stereo / structured light black-and-white depth cameras (Ensenso N35) are used. An RGB-D / stereo pair is mounted overhead looking down over the picking area, and the other stereo camera is mounted looking down the packing area, shown in Fig. 5.2. This setup allows us to combine both the color information from the RGB-D camera and the high-quality point clouds from the stereo cameras. The high-quality point cloud enables the pose recognition pipeline to achieve better accuracy compared to the RGB-D data.

The entire system layout is illustrated in Fig. 5.2. In the picking area, objects are placed

in arbitrary locations. The system assumes that objects are drawn from a known set, whose 3D models are available. In addition, all objects in the picking area will be packed, with their initial locations not touching. In the packing area, a predetermined box footprint area is marked off, and a box with matching dimensions is placed by hand to align with the footprint.

The software pipeline consists of calibration, perception, and planning. The calibration steps include:

- Factory-specified tool center point and camera intrinsic parameters were used.
- Camera extrinsic calibration using robot-mounted checkerboard
- The box footprint location is calibrated by jogging the robot to touch the footprint corner locations. The footprint corners are set to the tool center point coordinates relative to the robot base.

Object identification and pose estimation are then performed with the following steps:

1. Plane segmentation removes the tabletop from the point cloud.
2. Region growing segmentation [AB94] is used to segment object point clouds.
3. Pointcloud segments are projected onto the RGB image to obtain an object ROI.
4. The object is identified using a convolutional neural network model [HZRS15], trained on the experiment objects.
5. A CNN pose estimation model gives an initial estimate of object pose [WXZ⁺19].
6. If fitness [ZPK18] of the CNN estimation is below a threshold, point-cloud only template matching is performed.
7. The pose is fine-tuned with iterative closest points (ICP) [RL01b].

The packing algorithm [WH19b] chooses an object packing sequence and generates a grasp location and packing pose for each object. Top-down loading is performed. The planner checks several constraints: 1) no excessive force on the gripper caused by misalignment, 2) the loading trajectory is collision-free with the container position and objects already placed, and 3) stability of the object under gravity in the stack of existing items.

We suppose two forms of sensor feedback during/after packing execution with the robot arm:

- Force feedback, via the force/torque sensor on the UR5e’s wrist, is used for guarded moves during packing. This prevents crushing items.
- RGB-D and stereo cameras provide an overhead image of how items have been packed in the container.

5.3.2 Factors Affecting Packing Success

Packing success is defined by a final state having all objects unbroken and packed within the container’s dimensions.

Because the planner performs collision and stability checking, if the world behaved exactly as modeled, then the executed plan would be guaranteed to yield a successful packing. However, several errors are introduced by the execution pipeline.

We identify and describe nine key **error sources**:

1. Camera calibration error ϵ_{cal} .
2. Object identification error ϵ_{ID} .
3. Pose recognition error ϵ_{pose} .
4. 3D geometry modeling errors ϵ_{geom} .
5. Center of mass specification errors ϵ_{cm} .
6. Grasp positioning error ϵ_{grasp} .

7. Grasp acquisition (lifting) error ϵ_{lift} .
8. In-transit manipulation error ϵ_{manip} .
9. Box placement error ϵ_{box} .

Camera calibration error in the picking area causes deviations in the robot-relative object pose estimates, adding to problems acquiring a grasp. Camera calibration over the container leads to collisions with objects and walls of the container in the closed-loop packing strategy.

Object identification error leads the planner to choose a different 3D reference model than the actual object’s geometry. In packing setups where the correct subset of items must be picked from a larger collection, this could lead to more serious errors where the wrong item is packed. Our problem definition assumes that all presented objects must be packed, so object identification error is correctable during packing.

Pose recognition error leads to misalignment of the true object geometry and the 3D reference model used in grasp and packing planning. This causes grasp failures and inadvertent collisions during loading.

The 3D geometry of each object was acquired using a 3D scanner and multiple viewpoints. There is inevitable noise in these models, particularly noticeable in the underside of the object. Moreover, the scanned items were newly bought objects, and as items exhibited wear and tear during our experiments, some items underwent slight deformations. Although we do not have ground truth, we expect overall geometry errors to be less than 5 mm.

Centers of mass of objects are not measured precisely, and may also change with the object orientation (e.g., product settling in a box). This causes potential problems with grasp acquisition and inaccurate stability tests in the planner.

Grasp positioning errors occur when the suction gripper shifts the object during grasping. The related grasp acquisition errors occur when the gripper cannot achieve adequate suction to lift the object. In-transit manipulation errors occur when the robot prematurely

loses grip of the object, e.g., by causing large accelerations.

Finally, box placement errors can cause collisions with the sidewalls during packing.

Overall, these errors form cascading effects down the pipeline to contribute to unexpected events, such as toppling object placements or unexpected collisions, which lead to packing failures, as shown in Fig. 5.3. Informally, we refer to events in which an object hits a container wall or other objects prematurely as “smashing,” and events in which an object falls over or shifts as “toppling.” Note that the packing loop is performed N times, where N is the number of objects, so the system has an opportunity to measure and correct for errors accumulated during packing.

5.4 Error reduction methods

We propose two methods to reduce errors. The first strategy is closed-loop packing, which makes use of sensing feedback to adjust the packing plan so that if an object is grasped incorrectly or falls over during placement, subsequent steps can adapt to the new knowledge. The second strategy is robust planning, which predicts the future effects of uncertainty so that unexpected behaviors are less likely to occur. To evaluate the effectiveness, we perform experiments with four variations of planning and executing strategies.

1. **V1- Baseline:** Plans are generated and executed, assuming perfect models.
2. **V2- Robust planning:** The packing planner is modified to avoid tight fits.
3. **V3- Closed-loop vision:** The container sensor re-senses the object pile after each placement. If the heightmap of the pile deviates from the plan, then a replan is triggered.
4. **V4- Robust planning and closed-loop vision:** Both V2 and V3 strategies are employed.

Strategies V3 and V4 do not consider significant rearrangement of objects within the container (e.g., repacking), although we do allow the robot to push objects that have stuck

outside the container. We describe these strategies in more detail below.

We implement a grasping controller based on end effector force feedback, which is used in all experiments. During grasping, if the robot’s force sensor detected that the object was not successfully lifted, the robot will attempt the next computed grasp pose that can achieve the planned object placement with a collision-free loading motion. Moreover, while packing an object into the container, if a contact force exceeding 20N is sensed, the object will be released. This avoids crushing objects during “smash” events.

5.4.1 Open-loop packing baseline

This section provides details on the packing planner and vision system, which are used as-is in the baseline packing condition.

Packing planner

The packing planner algorithm from [WH19b] can be summarized as follows. The input is a set of N objects with reference geometries $\mathcal{G}_1, \dots, \mathcal{G}_N$, where $\mathcal{G}_i \subset \mathcal{R}^3$, and initial poses $\mathcal{P}^0 = (P_1^0, \dots, P_N^0)$, where $P_i^0 \in SE(3)$. Also, \mathcal{C} is given as the free space volume of the container. The problem is to find a placement sequence $S = (s_1, \dots, s_N)$ of $\{1, \dots, N\}$, grasp poses G_1, \dots, G_N , and placement poses $\mathcal{P} = (P_1, \dots, P_N)$, $P_i \in SE(3)$, such that each placement satisfies a set of robot-packable constraints.

Let $P_i \cdot \mathcal{G}_i$ denote the space occupied by item i when the geometry is at pose P_i . The solution packing should satisfy the non-overlap constraints

$$(P_i \cdot \mathcal{G}_i) \cap (P_j \cdot \mathcal{G}_j) = \emptyset, \forall i, j \in \{1, \dots, N\}, i \neq j, \quad (5.1)$$

containment constraints

$$P_i \cdot \mathcal{G}_i \subseteq \mathcal{C}, \forall i \in \{1, \dots, N\}, \quad (5.2)$$

for each $k = 1, \dots, N$, stability constraints

$$isStable(P_{s_k} \cdot \mathcal{G}_{s_k}, \mathcal{C}, P_{s_1} \cdot \mathcal{G}_{s_1}, \dots, P_{s_{k-1}} \cdot \mathcal{G}_{s_{k-1}}) \quad (5.3)$$

and manipulation feasibility constraints:

$$isManipFeasible(P_{s_k} \cdot G_{s_k}, P_{s_1} \cdot \mathcal{G}_{s_1}, \dots, P_{s_{k-1}} \cdot \mathcal{G}_{s_{k-1}}). \quad (5.4)$$

Non-overlap and containment constraints are calculated using standard collision detection methods. Stability is checked using a convex program, while manipulation feasibility requires 1) finding a feasible grasp pose and 2) ensuring that a top-down loading motion is collision-free both for the selected object as well as the robot.

For grasp planning, we use a point cloud-based planner that generates a set of vacuum grasp candidates, ensuring that the grasp location is nearly flat, the suction direction is nearly normal to the surface, and the center of mass is almost underneath the grasp point.

Specifically, the grasp point must lie within a radius $r = 2$ cm in the horizontal plane to the centroid of the observed point cloud segment. The areas under the gripper should be planar areas (80% of the surface points directly below the tool are within 0.3mm to the estimated plane), which helps the vacuum opening to grasp normal to the surface.

For use in the planner, the set of grasp candidates is tested one by one for manipulation feasibility, i.e., with the chosen grasp pose, the robot can transfer \mathcal{G}_i from P_i^0 to P_i without collision and with successful IK solutions being found along the trajectory. Moreover, in the packing pose, the resulting gripper axis needs to be within a tilting angle $\theta = \pi/4$ to the Z-axis to prevent excessive torques being applied at the gripper.

Pose recognition

The pose recognition pipeline estimates P_i^0 for $i \in \{1, \dots, N\}$ for use in the planner.

First, we acquire partial pointcloud segments for $\mathcal{G}_1, \dots, \mathcal{G}_N$. This is done through a process of plane segmentation and distance calculation that removes all observation not on the picking station. The remaining pointcloud is segmented by Point Cloud Library(PCL)’s [RC11] implementation of region growing segmentation [AB94]. The algorithm partitions the points that are close enough in terms of a smoothness constraint.

Next, the partitioned pointcloud segments are projected on the color image obtained with a RealSense camera. Pixels inside the bounding box of the projects are cropped

into a color image of each individual object. The ID of the object is classified with ResNet18 [HZRS15], trained on the experiment dataset using Pytorch [PGM⁺19].

For the experiment, we assume that $\forall i \in \{1, \dots, N\}, \mathcal{G}_i$ is sitting on a tabletop with pose P_i^0 , where the orientation component of P_i^0 is a *planar stable* orientation for \mathcal{G}_i . Therefore, pose recognition estimates the translation of \mathcal{G}_i while R_i^0 is from a finite subset modulo rotations about the vertical axis. We initialize the estimate R_i^0 from the Dense Fusion [WXZ⁺19] algorithm. Inputs to Dense Fusion include a pointcloud, cropped color image, and ID for the corresponding object segment. If fitness [ZPK18] for this estimation is below a certain threshold, a template matching technique is used as a fallback. Template object point clouds are generated at planar-stable orientations, and the best-fit match to the point cloud observation is selected.

Whether Dense Fusion or template matching is used for the initial alignment, we follow this with a dense point-to-plane Iterative Closest Point (ICP) [RL01b] alignment for fine-tuning. This is performed using the implementation from Open3D [ZPK18].

5.4.2 Closed-loop packing

In closed-loop packing, we monitor and remodeled the object pile after each placement. A failure is detected if the pile shifted significantly from expected, preventing subsequent items from being executed as planned. In case of failure, we replan the packing locations of the remaining items while maintaining consistency in the previous plan where possible. We also detect whether an object is lying partially outside of the container, and perform a *push maneuver* to attempt to push the object back in.

Specifically, with the overhead depth camera, we capture a pre-placement depth map of the container. If the incoming item’s 3D model at its planned packing location is in collision with the pre-placement heightmap, we count the number of colliding pixels. A “significant deviation” is triggered when this count exceeds a threshold of 20% of pixels belonging to the item.

If no significant deviation is determined, the item is placed to the planned transform as

usual, which helps maintain coherence with the existing plan. If a significant deviation is detected, the strategy replans a new optimal placement for the infeasible object given the pre-placement heightmap. Rather than trying to identify the pose of each object in the pile, which is prone to significant errors, we simply treat the heightmap as the geometry of an object rigidly-fixed to the container. The use of the sensed heightmap as a collision geometry suffices for collision checking between the robot and gripper, and stability checking still verifies whether there is sufficient friction between a rigid pile and subsequent placed objects. However, treating the pile as a rigid, unmovable object incurs some loss of information for the stability constraint checks, so the planner may decide to put the object in a location an underlying object may shift or topple.

To perform the push maneuver, we capture a post-placement heightmap, and detect whether any object’s horizontal extent surpasses the dimensions of the container. If so, we aim the robot’s gripper to push in a straight line perpendicularly to the box wall, with the end effector’s lowest portion slightly above the box’s maximum height. This movement is aimed at the geometry center of the portion of the object outside the container dimensions.

5.4.3 Robust planning

Robust planning avoids tight fits with a threshold distance of δ . Placements closer than δ to the container wall and placed objects will be penalized in the offline planner.

In [WH19b], a candidate placement is scored as follows:

$$c \cdot (X + Y) + \sum_{i=0}^{w-1} \sum_{j=0}^{h-1} H'[i, j] \quad (5.5)$$

Where X and Y are the translation of the object inside the container at a given orientation. H' is the sum of the container heightmap assuming the object is placed at the candidate placement. The heightmap is a 2D image of width w and height h , and c is a small constant. For robust planning, we add an extra term to the scoring function to

penalize tight fit:

$$c \cdot (X + Y) + \sum_{i=0}^{w-1} \sum_{j=0}^{h-1} H'[i, j] + C / \max(d_{min}, \delta) \quad (5.6)$$

Where d_{min} is the minimum horizontal distance between the item and the pile or the container wall. This is computed from the heightmaps of the pile and the item being placed. To compute the minimum distance, we only consider pixels in the pile/container heightmap whose pixel height is larger than object placement height Z and x-y location within and near the object 2D projection onto the x-y plane. The minimum of all horizontal distances between object pixels and the pile/container pixels is taken as d_{min} . The value C is taken to be a large constant such that C/δ dominates the other two terms in the sum.

It should be noted that conservative planning may fail to find a plan when non-conservative planning would succeed. To address these cases, the planner tries to replan with a smaller δ that decreases linearly until it is 0. If the $\delta = 0$ case fails, we declare planner failure.

5.4.4 Closed-loop packing and robust planning

In our final experimental strategy, we combine both closed-loop repacking and robust planning. This is mostly a straightforward combination, except that we note that when δ is decreased due to a planning failure, we maintain the decreased value for subsequent planning steps.

5.5 Analysis and experimentation

A systematical evaluation is performed on the proposed packing system that studies the sources of error and model its magnitude. Experiments are conducted to quantify the impact of such errors on the overall packing success rate in Monte Carlo simulation with bullet simulator [CB19] and on the physical testbed.

5.5.1 Itemset

15 real-world items from the YCB video object dataset [XSNF17] are selected that are mostly rigid, opaque, graspable, visible to the depth camera, and in a wide range of shapes and weights (Fig. 5.6). High-quality object meshes for the 15 items are available so that we can simultaneously perform the evaluation on both physical and simulation testbeds.

5.5.2 Uncertainty evaluation

We collected a dataset of all experiment items placed in various poses and positions on the picking area, complete with RGB images, depth images, and pointclouds, taken by Realsense and Ensenso cameras in their overhead mounting positions. Around 200 poses and positions are taken for each object. 70% of the dataset was used for training, and the remaining 30% of the dataset was withheld for testing.

Object identification

In our test set, object segmentation was 100% successful because the input objects were placed in isolation. Our object recognition pipeline achieved a 98.5% top-1 accuracy on the testing set, which indicates that for 7.2% of 5-item order sets, at least one of the objects will be misidentified. Note that because all objects in the packing area need to be packed, a misidentified object will not necessarily lead to failed packing. Instead, it contributes to a (potentially substantial) geometry modeling error.

6D pose estimation

To obtain the pose estimation error, the 6D pose estimation result from our pipeline is further corrected by manual alignment of the projected and ground truth (using the Meshlab software). We considered using an alternative approach to pre-define a ground truth pose and manually align the object to this pose, but we found that symmetric objects (e.g., boxes, cylinders) induced large pose error measurements, even though these errors would

have virtually no effect on picking and packing. On our test set, the average translation and rotation error is 0.7 mm and 0.3 rad.

Rare errors

Grasp acquisition errors were not observed in our setup (so $\epsilon_{lift} \approx 0$), in large part because we chose objects from the YCB dataset that were more easily lifted by suction. In-transit manipulation errors never occurred during our experiments ($\epsilon_{manip} \approx 0$). Box placement error is also quite low ($\epsilon_{box} \approx 0$) compared to other sources of error, e.g., calibration, because it is relatively simple to align the box to within a millimeter or two of the calibrated footprint. Moreover, it would be a simple matter to add rigid fixtures to the packing area to constrain the box even further. Since these errors are so low, we do not include them in our analysis.

5.5.3 Summary statistic: placement error

Observe that calibration, pose estimation, and grasp positioning errors all contribute to the pose of the object-in-hand, which leads to an ultimate error in where the object is packed. Specifically, this is the final translation and rotation error of an object from a planned packing location P_i , assuming no collisions occurred during loading. We call this summary error *placement error* ϵ_{place} , and can be evaluated experimentally. The evaluation performs 80 single item manipulations by a known transform that transforms object from one planar orientation to another. Before the manipulation, the pointcloud of the object is captured and the pose of the segment is estimated. After manipulation, the pointcloud of the object is captured again, and the observation is manually aligned to the projected object pose. ϵ_{place} is taken to be the RMSE of these alignment errors. Our experiments show that placement error is on average 1.02 cm in translation and 0.41 rad in rotation.

A summary of all error sources is given in Tab. 5.1.

Table 5.1: Summary of empirically evaluated error sources

Error source	Magnitude
ϵ_{ID}	1.5%
ϵ_{pose}	0.7 cm (translation), 0.3 rad (orientation)
ϵ_{geom}	Estimate \approx 0.5 mm
ϵ_{place}	1.02 cm (translation), 0.41 rad (orientation)
ϵ_{lift}	\approx 0
ϵ_{manip}	\approx 0
ϵ_{box}	\approx 0

5.5.4 Simulation testing

A more extensive set of experiments were conducted in simulation, where it is more practical to systematically introduce different magnitudes of errors and larger test sets. Varying degrees of error from ground truth object poses and depth sensor readings are used as input to each of our packing strategies. The simulations are run on several desktop computers with 32GB of RAM and Intel I7 processors.

For these tests, we randomly generated 1000 sets of 10-item orders from YCB [CSB⁺17] and the APC 2015 object set [Rut15]. For consistency, we use the same item sets across the range of packing strategies. For picking, we rigidly attach objects to the robot’s gripper with a randomly-generated pose error, using the same standard deviations as ϵ_{place} , but scaled with an error scaling factor from 0-200%. The object is dropped when the robot reaches the desired placement pose, or a 20N force threshold is reached.

The packing success rates for strategies V1-V4 are plotted in Fig. 5.7. For V2 and V4, a δ value of 1 cm is used. At 0% error, all experimental variations performed similarly, just under 98% success rate, while closed-loop strategy performs slightly worse. This can be largely attributed to the treatment of the pile as a rigid, infinite-mass object during

replanning, which fails to capture the nuances of the pile’s stability.

When an error greater than 100% is introduced, the success rate for the baseline strategy drops off sharply as the error increases. Meanwhile, the rate of drop off is noticeably slower for the other 3 comparison strategies that employ robustness measures. The best performer is V4 that uses both robust plannings as well as dynamic replanning, achieving the highest or near highest success rate at all positioning errors introduced, and maintain a packing success rate at 94.5% at 2.04 cm translation error and 0.82 rad in rotation error.

It should be noted that increasing δ does not always lead to better performance under uncertainties. A δ of 2 cm was also tested in V2 and V4, achieving similar and slightly worse performance than a δ of 2 cm. With a higher margin, more planning under margin will fail and fall back to use a smaller margin.

5.5.5 testing on physical platform

For testing on the physical platform, we generated 100 sets of 5-item orders from our experiment items. The robust planning strategies (V2 and V4) use $\delta = 1.02$ cm. Tab. 5.2 summarizes the results. Both closed-loop and robust methods (V2 and V3) perform significantly better than the baseline, while the combined method (V4) performs best of all, beating the baseline success rate 98% to 83%.

We also show the number of times the maximum loading force (20 N) was exceeded, which indicates approximately how often a smash event was observed and how much damage the object sustained during packing. The same pattern is observed, showing that both closed-loop and robust strategies decrease the rate of smashing. The last row shows how many times a closed-loop push maneuver was performed, showing that these are rather rare, but are still helpful to improve success rates by a few percentage points.

We recorded object classification results, pose estimation results, as well as planned and final placement for each experiments, which allows us to examine causes of cascading failures more closely. In V1, 9 failures occurred with one or more items protruding outside the box footprint, while 7 failed for objects within the container but exceeding the maximum

Table 5.2: Success rates on the physical packing platform, 5-item orders

	V1	V2	V3	V4
Success rate (%)	83	95	95	98
Max force exceeded (# per order)	2.22	1.40	1.12	0.7
Push performed (# per order)	—	—	0.07	0.03

height. A final case failed when a rigid object tipped over and stood up underneath the gripper, causing the robot to report a “maximum load exceeded” fault. Some failure cases can be traced to one predominant factor. 4 failures can be traced back to failed object identification, which leads to wrong shape estimation of the object to pack. 3 cases can be traced back to significant pose estimation error (e.g., rotation of 90 degrees, etc.). 1 case can be traced back to object tipping, which led to sudden force changes too rapid to be caught by the force control loop. The remaining failure cases are caused by cascading errors, e.g., slightly misplaced objects caused the object underneath to tip over, resulting in the subsequent objects to be packed exceeding the container height.

In V2, 3 cases failed when one or more item protruding outside the box footprint, and the other 2 failed during toppling events, causing items to exceed the max height. Among the failure cases, 2 can be traced back to misidentified objects, while 1 can be traced back to significant pose estimation error.

In V3, all 5 cases failed due to exceeding the max height. All objects protruding outside the box footprint were corrected by pushing maneuvers. Among the failure cases, 2 can be traced back to misidentified objects, and 2 can be traced back to significant pose estimation error.

In V4, the 2 failure cases appear similar – the *cracker box* item toppled and exceeded the max height. One case was caused by a significant pose estimation error, while in the other, the *bleach* item was pushed onto the *cracker box*, and the impact caused it to tip

over.



Figure 5.1: The experimental packing setup consists of a UR5 robot equipped with E-pick suction gripper, and overhead cameras (not pictured). Errors from pose recognition, camera-robot calibration, box location calibration, grasp planning, the center of mass estimation, and geometric modeling affect the overall success rate of packing.




Figure 5.2: Top-down diagram of the packing experiment. Objects to be packed are arranged in the picking area, and must be packed in the container (e.g., a shipping box). Overhead cameras (Realsense SR300 and Ensenso N35) provide color and depth information.

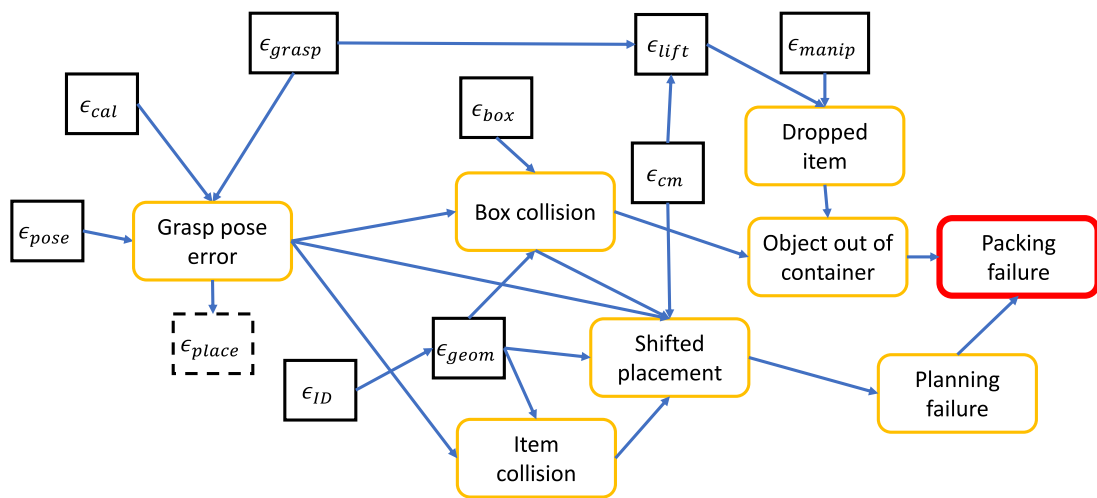


Figure 5.3: A diagram of how error sources (rectangles) contribute to unexpected events (rounded rectangles) and ultimately to packing failure. We use ϵ_{place} as a measure for grasp pose error.

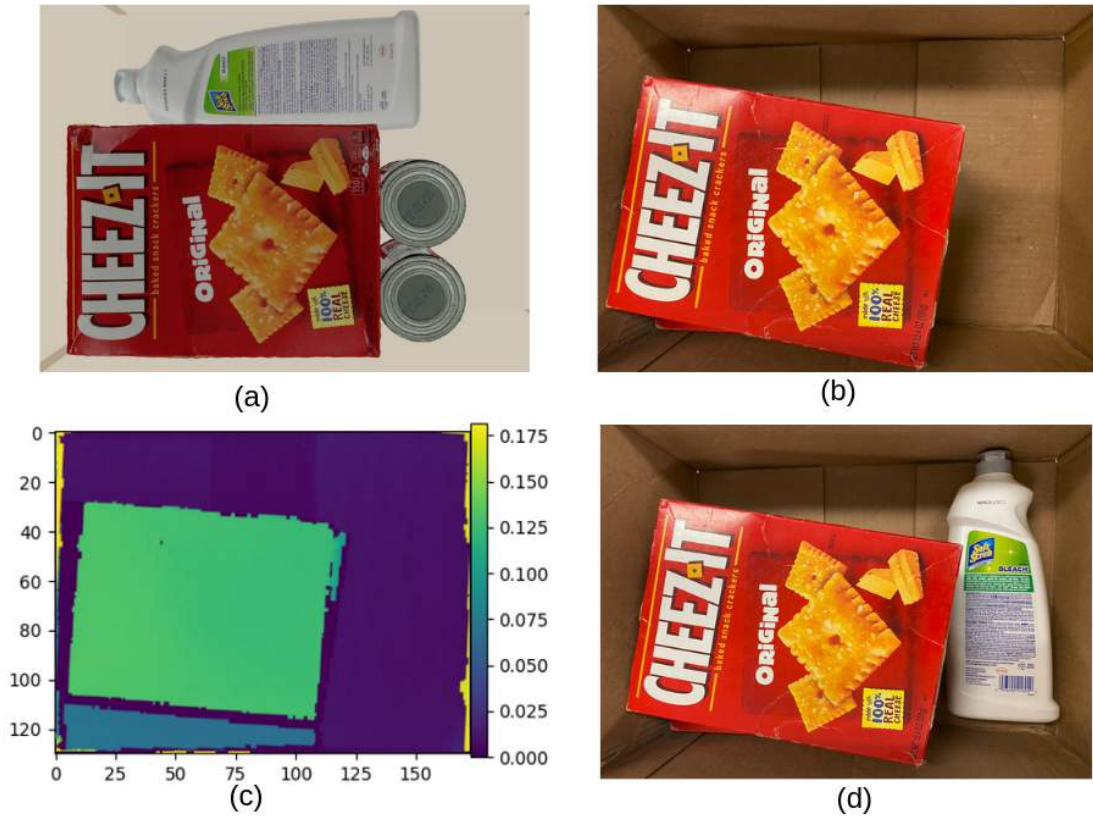


Figure 5.4: Closing the loop around vision leads to more robust packing. (a) Offline plan. The 2 *cracker boxes* are stacked first, followed by *bleach* and 2 *soup cans*. (b) The second *cracker box* is shifted during packing, blocking the access for *bleach*. Open-loop packing would lead to a failure. (c) The closed-loop strategy captures a heightmap of the placed pile, detects a violation for the *bleach* placement, and plans a new location. (4) *Bleach* is placed successfully in the empty location.



Figure 5.5: Conservative geometry margins reduce inadvertent collisions. (a) A non-conservative plan (left) causes a failure during the execution of the plan (right). Specifically, the shifted *cracker box* caused a failure and crash of the *sugar* item planned on top of it. (b) A conservative plan with $\delta = 1$ cm margin. During execution, all items are placed within the container without causing an inadvertent collision.



Figure 5.6: 15 items from YCB video object dataset are used in the packing experiment

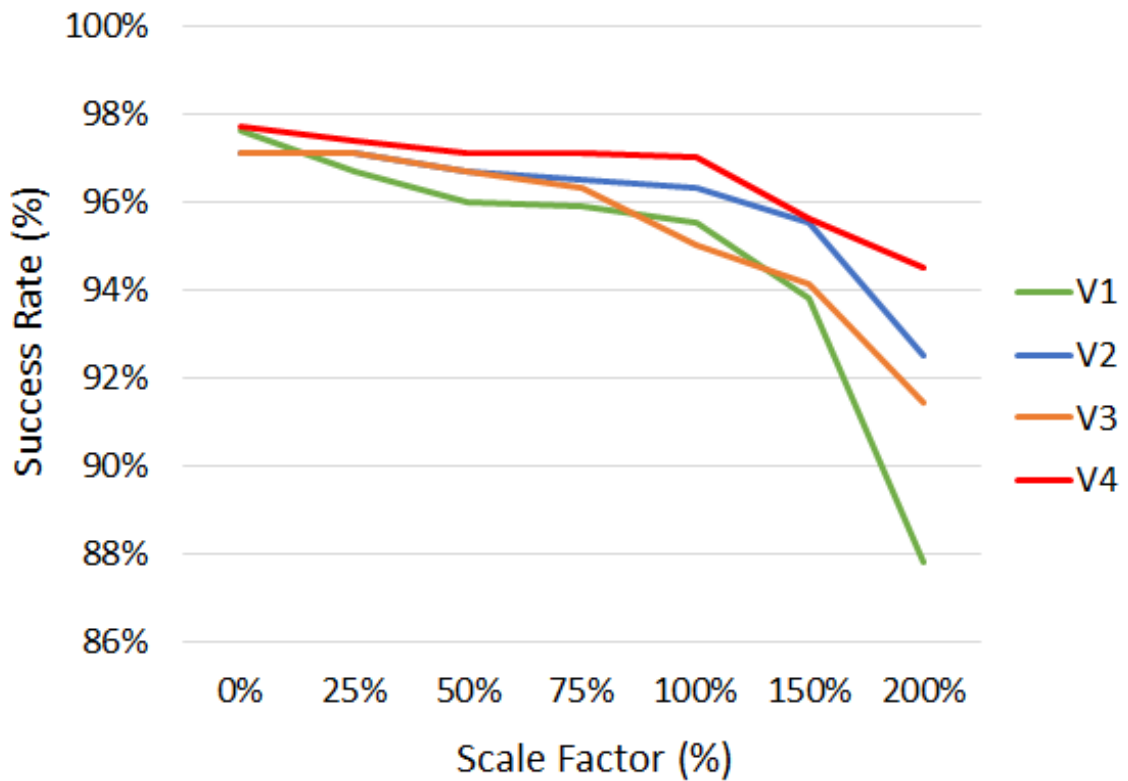


Figure 5.7: Simulation packing success rates for 10-item orders. Pose error scaling factor is given on the horizontal axis, and the success rate is given on the vertical axis. The baseline (V1) success rate drops off dramatically as pose error increases, while the robust technique (V4) is much less sensitive to error.



Figure 5.8: Example of failed packing executions, for each experimental condition. In the baseline (V1), all items are planned tightly, and the *cracker box* hit a container wall, causing it to tilt, and for all subsequent items planned on top of it to smash. In robust planning (V2), although margins were enforced in planning, a substantially incorrect pose estimation caused the *bleach* item to stick out of the container, and no action is taken to correct it. In closed-loop packing (V3), the first-placed *cracker box* caught on the side and failed to be pushed in. This was the only case when the push maneuver did not work in our examples. Although this caused a failed plan, replanning prevented the subsequent items from smashing onto the *cracker box*. With both closed-loop packing and robust planning (V4), a replan was triggered after the *cracker box* and before the *bleach* item. Because the pile was treated as a fixed, rigid object, the *bleach* was placed on top of the *cracker box*, but since it was unstable, this caused it to tip over.

Chapter 6

Robot Button Pressing In Human Environments

In order to conduct many desirable functions, service robots will need to actuate buttons and switches that are designed for humans. This paper presents the design of a robot named SwitchIt that is small, relatively inexpensive, easily mounted on a mobile robot, and actuates buttons reliably. Its operating characteristics were developed after conducting a systematic study of buttons and switches in human environments. From this study, we develop a categorization of buttons based on a set of physical properties relevant for robots to operate them. After a human calibrates and annotates buttons in the robot’s environment using a hand-held tablet, the system automatically recognizes, pushes, and detects the state of a variety of buttons. Empirical tests demonstrate that the system succeeds in operating 95.7% of 234 total buttons/switches in an office building and a household environment.¹

¹This chapter is reproduced from Fan Wang, Gerry Chen, and Kris Hauser, “Robot Button Pressing in Human Environments,” in 2018 International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 2019, May 2018, pp. 7173–7180.

6.1 INTRODUCTION

Recent years have seen a growing interest in service robots that assist humans in their daily lives, such as in households, offices, factories, and hospitals. Interacting with physical switches and buttons is a pervasive part of human life, used to operate lighting, appliances, computers, elevators, and machinery, and will therefore be an important capability for these robots. As a result, these control devices have been designed for simple, intuitive, and reliable operation by humans, both in terms of their ergonomic mechanical properties and distinctive physical appearance. There are many types of control devices in human environments, including push buttons, toggles, slides, and knobs, and in this paper these devices will hereafter be considered synonymous to a *button* or a *switch*. Manipulating a button/switch to perform a desired effect may also be variously referred to hereafter as *button pressing*, *switch operation*, or *switching*. Operating switches with a robot with human-level ease and reliability remains a challenging task, due to the fundamentally different sensing and actuation modalities on robots vs. humans.

This paper presents the system development and design of a compact autonomous button operating robot called SwitchIt. It is a 3DOF device based on relatively inexpensive sensing and actuator hardware (Fig. 6.1). A short manual calibration setup is performed once for a given environment using a handheld tablet and fiducial markers to identify the identity and purpose of each button. After calibration, the system recognizes and operates a wide variety of buttons automatically. It can also use sensor feedback to detect the state of many buttons and whether they have been successfully operated.

The system integrates three primary contributions:

1. We perform a systematic categorization of over 600 buttons and switches found in offices and homes into 6 classes based on their physical properties required for robotic actuation. We propose a taxonomy of buttons from a robot’s operational point of view, and characterize several relevant physical properties of these buttons, including travel, size, shape, and operating force.




Figure 6.1: SwitchIt is a spherical robot equipped with an RGB-D camera. Button panels are annotated using a QR code sticker affixed during a manual setup phase. Shown here mounted on a tripod, the robot is preparing to press buttons on an electronic passcode panel.

2. We develop an annotation, calibration, and perception subsystem that achieves high-reliability button recognition, localization, and state detection. The calibration process also handles reflective and dark surfaces.
3. We design a compliant, scalloped end effector tip that can actuate pull buttons and turn knobs, and is robust to positioning error.

In controlled testing, our perception system localizes buttons with < 2 mm error and detects the state of toggle switches with 100% accuracy. For typical localization errors, the scalloped

tip design achieves 99% repeatability compared with 89% for a cylindrical tip. We also test the system in an uncontrolled office and home environment, with 234 switches attempted in total. The platform succeeds at operating 224 (95.7%) total switches. In particular it was highly reliable at operating push buttons, sliders, rockers, and switches.

6.2 RELATED WORK

Computer vision techniques have long been employed for service robot to navigate and interact with objects in human environments [MWG⁺10, PLS13]. Specifically, several authors have addressed button identification issues that uses visual features in RGB images to identify and locates buttons. Identifying features could either come from a priori knowledge of the type of button [WHLC10, KAO07] or based on results of machine learning [KCRN10]. Most prior research using features based on prior knowledge is performed on elevators buttons [WHLC10, KCRN10, KAO07]. The advantage of this approach both seen and unseen buttons can be detected and to some degree understood automatically. However, these algorithms usually must rely on contextual cues such as grid layout and sequential arrangement of the buttons. Machine learning strategies are also commonly used [SUB96, BKP11]. Sukhoy and Stoychev (2010) use an active machine learning strategy to identify and trigger a button autonomously. This method was used to train a robot to identify the active part of a door bell, and to trigger it effectively [SS10]. The auditory feedback of the feedback is used to determine whether the button was successfully pressed.

Several authors have addressed button identification issues with autonomous recognition algorithms that use visual features in RGB images to identify and locate buttons. Identifying features could either come from a priori knowledge of the type of button [WHLC10, KAO07] or based on the results of machine learning [KCRN10]. Most prior research using features based on prior knowledge has been performed on elevator buttons [WHLC10, KCRN10, KAO07]. The advantage of this approach is that both seen and unseen buttons can be detected and, to some degree, understood automatically. However,

these algorithms must usually rely on contextual cues such as a grid layout and sequential arrangement of the buttons, or be applied to only a limited class of buttons.

Our approach asks for a small amount of environment augmentation and manual labeling to identify each button definitively. Environment argumentation has been used in other robotics systems as well to aid in object identifications. This removes ambiguity and achieves much higher accuracy than automated identification while require minimal setup time. Tools most commonly used are augmented reality tags such as RFID tags [NDRK09], QR code [LLTK14] and other artificial marks [KOM⁺03, KOT⁺03]. Those tags usually provide information on the location of the objects, instructions on how to interact with this specific object and a task completion criteria. We use a similar approach with QR codes, which was also previously used to enable a mobile manipulator to plug itself in [LLTK14].

The work arguably closest to ours is Nguyen et al (2009). They uses a combination of an augmented environment and a variety of sensors to help a robot interact with its environment [NDRK09]. This work does allow the robot to operate certain light switches. Force sensing and visual feedback, in the form of a change in lighting condition, is used to detect the change in button condition. In our work,we further demonstrate that high accuracy localization can be achieved using AR marker and RGB-D sensor alone, we also propose parameterized motion primitives for each class of buttons as opposed to the explicit defining instructions for each individual object to interact with.

6.3 CHARACTERIZATION OF SWITCHES IN HUMAN ENVIRONMENTS

A button or switch changes its internal electrical connection or signal based on the force applied to its external active mechanical component. For a human or robot operator, the underlying circuitry of a switch can be considered a black box and can be mostly ignored. The main focus is performing the appropriate physical action to correctly and safely trigger the switch.

Although some industrial settings employ switches that require significant force (or even tools) to be applied, here we focus on switches that are designed to be operated by one or two human fingers. These switches are designed with size, shape, material, and mechanical resistance that are comfortable for human fingers to manipulate. Moreover, a switch usually has a distinctive appearance that indicates its mechanical functionality and semantic meaning, and usually provides feedback that can be promptly perceived and interpreted.

Although switches are ergonomically designed, the movement needed to trigger a switch safely and reliably is actually a delicate skill, acquired by humans through years of practice. Humans use memory, visual feedback, tactile feedback, and a variety of finger and hand contact strategies, and also progressively improve the efficiency and comfort of switching motions. For example, to operate stiff switches, a senior citizen with reduced hand strength will adopt finger postures that apply more leverage to stiff switches.

We collected data for over 600 switches in office and home environments to help design our robots operational characteristics. This section describes their typical distribution and operational characteristics.

6.3.1 Button Taxonomy

Laypeople usually address buttons by a common name that references its function, such as light switch, toggle button, touch pad, dimmer, or keyboard. Electricians categorized them on the basis of their electrical connection such as single-pole single-throw (SPST) or double-pole single-throw (DPST), or by their triggering mechanism (sliders, push buttons) or the type of the application it is used on (light switches, dimmers). For a robot, perhaps the most useful categorization of switches is in terms of its physical triggering mechanism.

The operating mechanism of buttons and switches can be described with respect to a normal direction facing outward from a *button panel* (a plane behind which the electrical circuits are hidden). Our proposed taxonomy divides all household buttons into following 6 types based on their operating mechanism (Fig. 6.2).

- **Push-buttons** operate via application of a force that moves the operating part inward toward the panel. The movement of the button is linear.
- **Toggles**, such as in household light switches, generally have a rod-shaped protrusion (known as a *level*) that can be rotated about an axis to toggle between 2 distinct states. Some toggle switches have 3 or even more states. Internally, a spring and plunger mechanism is used to aid in operation, and equilibrium is only achieved in the extreme positions.
- **Sliders** require lateral movement, but the level can remain in equilibrium in any state in its travel range. Typically a slider is operated with a linear motion and stays in place using friction. However, some have discrete equilibrium states enforced by internal springs.
- **Rockers** are pressed on one of 2 ends like a seesaw and toggle between two discrete states. Unlike a toggle, a rocker is triggered by applying an inward force, primarily normal to the button panel.
- **Turn knobs** rotate along a center axis perpendicular to the panel face to adjust either a continuous or discrete value. Humans typically operate knobs using much of 2 or more fingers to achieve sufficient tangential friction about the axis.
- **Pull buttons** operate by a pulling action that moves the operating part away from the base to open or close the contacts. To operate the button, it must be solidly grasped either on the back or via friction on the sides.

6.3.2 Operation

Switches operate in three typical patterns: Momentary, Alternating, and Latching. In momentary operation, the switch is in an active state only while a force is being applied, and then once unloaded, a spring returns the switch to its original position. For alternating operation, the switch's state is held after it is released, and a different force must be applied




Figure 6.2: Buttons characterized by type.

to change state (this characterizes almost all toggles, pull buttons, and rockers). In latching operation, a second force in the same location/direction of the initial force returns the switch to its initial position, which is accomplished by some spring-loaded toggling mechanism. Latching characterizes several types of push-buttons.

6.3.3 Location and Geometry

The location of the button in the environment, direction of travel, and the travel distance affect whether the operational capabilities are within the workspace of the robot. We are primarily concerned with height, but to ensure accessibility of a mobile base it may also be important for a robot designer to consider surrounding obstacles and clutter. The geometry of the button and its relationship to surrounding buttons is also an important aspect of finger design, since it is important to be able to press the button without accidentally activating nearby buttons.

In our survey, the vast majority of buttons on the walls or doors have a height of 1.06–1.44m above the ground. The exception is elevator buttons, which have height range 0.88–1.72m. (The highest button is designed to be difficult to reach; it is to be used only in case of emergency.)

6.3.4 Force

There are different types of forces associated with button pressing, but we primarily focused on Operating Force (OF) which is the peak force needed to change the state of the button. The typical force-stroke characteristic of a button displays no movement until a breakaway force is reached, after which the force increases with increasing displacement until the peak at the operating point. Afterwards, the force follows a sharp decrease, and then gradually increases again until reaching the total travel distance [OMR17]. If tactile sensing is available, force profiles can be very helpful for detecting the success of button pressing. However, our system uses encoder derivatives to estimate the applied force.

Table 6.1: Breakdown of button characteristics by type

	Push	Toggle	Slider	Rocker	Turn	Pull
Prevalence	66%	25%	2%	2%	2%	4%
Force (N)	0.7-12.5	2.6-6.2	0.3-0.7	2.4-9	0.7-15	3-20
$\overline{\text{Force}}$ (N)	7.15	4.61	0.5	4.5	4.2	10.2
Trav. (cm)	0.1-0.35	1-2	0.5-11	0.2-0.45	1-5.2	0.2-0.75
$\overline{\text{Trav.}}$ (cm)	0.167	1.53	4.3	0.3	2.4	0.4
$\overline{\text{Sep.}}$ (cm)	1.6	8.5	X	X	X	X
$\overline{\text{Height}}$ (cm)	119	126	122	118	76	132

Tab. 6.1 gives results of our survey, listing range and mean value of operating force (Force and $\overline{\text{Force}}$, respectively). Most buttons have OF in the 4.5–8.5N. 19.7% of buttons have OF > 8.5N and 20.8% have OF < 4.5N. Only two buttons exceeded OF > 12N. For example, an elevator emergency stop pull button required 22N to operate. However, we find these buttons are not designed to be used on a daily basis.

It should be noted that all buttons we surveyed have been in operation for at least one year. New buttons are usually much harder to activate and require some usage before their internal springs soften.

6.3.5 Surface material

The body of electronic switches are usually made with metal or plastic. While highly polished plastic and metal makes attractive appearances, reflective surfaces are challenging for depth estimation with sensors. Our survey shows that 39% of buttons are made with highly polished plastic or metal and 10% are made with black material.

6.3.6 Travel distance

There are two parameters regarding travel distance: the total travel (TT) and operating travel (OT). TT indicates distance to a hard stop, while OT indicates distance until the switch is triggered. We are primarily interested in determining OT, although TT may be useful for tactile feedback. Most switch specifications suggested TT-OT to be between 0.5mm-1.8mm [C&17, OMR12].

In our survey, OT varies significantly, particularly among sliders (Tab. 6.1). However, we are most interested with the travel of push buttons, since they provide the least reliable visual cue of button state. From this data, we determined that an open-loop position controlled robot could press down a maximum of 5mm before linear actuator motor stop is detected to maximize its chance of activating a button successfully, while also being unlikely to cause damage to the button.

6.4 THE SWITCHIT PLATFORM

With the above data in mind, the SwitchIt robot accessory is designed to operate a large number of switches and to be easily mounted on a mobile robot platform. Setting up the system for use in a new environment requires a human to first perform a calibration procedure, which involves affixing QR codes to button panels and annotating reference models of the panels using a hand-held tablet. Afterward, the system will autonomously recognize any visible panel, suggest a reference position for the robot's base, and once in position, press a requested button or sequence of buttons.

6.4.1 Hardware

The robot arm used in our system is a custom 3DOF spherical robot that can pan, tilt and extend. The pan-tilt DOFs are ScorpionX MX-64 Robot Turret kit item number KIT-SXT-MX64 and extension is provided with a 50mm Firgelli Linear actuator. We have tested the physical capability of the arm in its workspace and measured a output force exerted

by the tip in a range from 6N to 18N. Although the robot is generally weaker and more susceptible to flexing the further it moves away from the center of the workspace, it should nonetheless be strong enough to trigger most buttons, which have operating force $< 12\text{N}$. The positioning accuracy of the robot after calibration is measured to be sub-millimeter on average, and less than 3mm maximum.

For sensing we use a single Intel RealSense F200 RGB-D camera to do colour and depth capture. These cameras are inexpensive (purchase price \$129), have a depth range of 20–120 cm, and work optimally in well-lit, indoor environments and diffuse objects.

Both the robot and camera are mounted on a fixed base using Plexiglass and a camera mount. We assume that the mount is attached to a mobile robot or an arm that has sufficient rigidity to keep the unit roughly in place while it presses buttons.

6.4.2 Scallop fingertip

We considered using a cylindrical rubber-coated tip, with similar shape and size to a human-finger. However, a novel scallop design proved to manipulate buttons much more reliably. The principle of the design is to increase tolerance against positioning error, and to enable motion primitives of pulling and turning which are usually difficult to achieve with a 3-DOF robot.

The design contains four rigid scallops protruding from its lateral edges (Fig. 6.3). For switches with activating rods, a scalloped channel guides motion towards the center-line, which corrects for positioning error and increases effective lateral friction. An underlying skeleton is 3D printed from rigid plastic, and this is coated with a 1-3mm thick PMC-121/30 rubber compound. The rubber coating provides compliance and large friction that helps correct for positioning errors and reduce slippage. The coating at the forward tip of the finger is curved with a dimension and shape similar to a human’s index finger. The scallop protrusions are also coated with longer rubber “skirts” to establish larger contact area when turning knobs. Finally, the rigid skeleton is also designed with a narrow “hook” located on one side of the finger, which is designed for holding onto pull buttons.




Figure 6.3: Cylindrical and scallop tip designs: (a) Scallop tip skeleton, (b) Scallop tip with coating, (c) Side view of tips. From top to bottom: scallop skeleton, cylindrical, scallop coated with rubber, (d) Top view of tips. From top to bottom: scallop skeleton, cylindrical, scallop coated with rubber

6.4.3 Environmental annotation and calibration

To apply the method to a new environment, a manual setup procedure must be performed to populate a database of known button panels. The process is relatively fast and the environment is minimally altered. For example, to complete annotation, calibration, and information entry for a medical device panel with 10 buttons takes less than 4 minutes (Fig. 6.4). The procedure consists of the following steps:

- The user affixes a QR code on or near the button panel.
- Using a tablet with attached RGB-D camera, the user takes a picture of the panel and provides an identifier for the panel.
- Guided by the annotation GUI, the user adds each button by name, type, and designates areas of interest on the picture.
- The panel identifier, QR code, RGB-D information, button names, type, location, size, and areas of interest are saved to a database.

It should be noted that dark or highly reflective panels and buttons cause problems with depth estimation, which could lead to erroneous 3D button location estimates. To accommodate these types of materials, the user should temporarily apply matte tape (masking or painter's tape) to the button panel when capturing the reference RGB-D image. After calibration the tape may be removed.

One omission of the current procedure is that we do not store a 3D map of button panel locations. As a result, to use our system, a mobile base must be able to first position the camera to observe the panel's QR code. Future iterations of our system might record panel location, and incorporate simultaneous localization and mapping (SLAM) software to guide a mobile base to a desired panel.




Figure 6.4: Steps of the panel calibration process: (a) Taking a picture of the panel panel with RGB-D camera, (b) Tap on a button to zoom in. Already calibrated buttons will be marked, (c) Draw rectangles in the area of interest as guided for each type of buttons, (d) Zoomed in button details make it easier for operators to calibrate with higher accuracy




Figure 6.5: Point cloud and RGB view showing localization of button surface in real-time. Computed locations of the buttons are drawn as red rectangles.

6.4.4 Button Panel Recognition and Localization

Recognition and localization consists of an imprecise QR code localization followed by a more accurate point cloud registration via Iterative Closest Points (ICP) algorithm [BM92, CM92]. When a QR code is detected, the panel reference RGB-D image and all button annotations become available. A first guess is obtained from the QR code, which gives an estimated affine transformation between the reference image coordinates and the current camera coordinates. Since QR codes are relatively small, this estimate is often inaccurate.

To improve accuracy, we then apply ICP to match the point cloud corresponding to the reference RGB-D image to the currently observed point cloud. The QR code localization gives a reasonable initial guess for this optimization. It should be noted when depth data is missing or corrupted by dark or highly reflective surfaces, ICP is not as effective, and localization relies more on the QR code and surrounding non-reflective surfaces.

Although RGB-D camera is relatively cheap, one significant drawback of this kind of camera is that they are unable to do depth measurements on objects with black or highly reflective surface. Therefore people turn to lasers and other more expensive sensing equipments for robust sensing. Buttons and buttons panels are often painted in black or are made in very shiny material such as stainless steel. However, since we are only tracking the QR locations in real time and the current position of the buttons are calculated from




Figure 6.6: Calibration areas by button type: (a) Push button, (b) Toggle, (c) Rocker, (d) Slider, (e) Turn knob, (f) Pull button.

the location in the database transformed by the affine transformation, the program is not actually detecting the real time position of the buttons with the RGB-D camera so that it is no issue that the RGB-D is incapable of doing the depth measurement on those buttons. However, if the entire button panel and it's surroundings are invisible to the depth camera, our program don't have the enough point cloud data to perform ICP and need to rely on the transformation provided by the QR code only.

6.4.5 Button State Recognition

Many buttons provide tactile and/or auditory feedback primarily intended for humans, which can also be used by the robot to determine whether it has successfully been switched. For robot not equipped with microphone or force sensor, it can be quite difficult to determine the state of the button and to confirm the completion of the task. Some of the previous

research uses heuristics such as dimming in the lighting condition to determine if a light switch has been successfully switched off [NDRK09]. However, those heuristics are not always available or can be relied on.

SwitchIt uses RGB-D information to detect the state of toggle, slider, and rocker switches, as well as push-buttons with back-lit LEDs. The shape of different switch positions can be quite distinctive. The lever rod of a toggle switch rests on 2 opposite sides at different states, and rocker surfaces tilt up at different sides. Take a light switch as an example. If the switch is in the up position, the average distance from the button surface to the underlying plane will be greater on the upper side of the button.

Using a region of interest from the calibration data, our method calculates the average displacement from the panel plane in both halves of the switch, and detects the switch position by the maximum displacement. We filter out noise at 5 cm distance from the panel since this is most likely caused by obstacles in front of the camera, e.g., the robot itself. We use the same method for detecting states of rocker buttons and sliders.

State detection is challenging for push-buttons, since many do not change in appearance and shape after activation. Some push buttons do provide visual cues, for example, an LED back-light. We therefore focused on detecting these differences in the RGB image. However, we find that very few push buttons provide visual cues, and the interpretation is not always consistent (e.g. a backlight turns on vs the button itself lights up).

6.4.6 Control

The controller of the robot is initiated when the robot is in reach of the button panel, and a button pressing sequence has been specified. The robot performs one or more guarded end effector moves in Cartesian space, determined by the button type and areas of interest collected during calibration. The pushing strategy for each type of button is as follows:

- Push: Approach the center of the marked zone and push down. Pushing stops if one of the following conditions have been met: 1) the linear actuator has fully extended

by 5mm, or 2) encoder readings indicate that the linear actuator has been stopped for 0.2s.

- Toggle: Linearly interpolate between center of 2 zones, in the direction needed to switch off / on.
- Rocker: Linearly interpolate from zones 1 to 2 or from 3 to 4, depending on the operational state. (We found this diagonal movement to slip less frequently than pushing straight downward.)
- Slider: Same as the toggle, except that the user / supervisor can specify a fractional travel amount.
- Turn Knob: Button center locations, button radius, and knob depth are determined from the marked region. During actuation, the tip touches the side of the knob and moves in a circular motion in the direction specified.
- Pull: First, approach the side of the button with the “hook” pointing inward and then move inward by 5mm. Interpolate toward a point 5mm in front of the pull button surface center, or until the encoder reading indicates that the linear actuator motor has been stopped for 0.2s.

6.5 EXPERIMENTS

We have done an full accuracy measurement of our system and separate localization errors with and without ICP. We also test our system with test panels that contains buttons of various types, shapes and stiffness, and finally an exhaustive real-world test in office and home environments.

6.5.1 Calibration

The robot and camera are automatically calibrated using a colored cross shape fiducial fixed to the end of the robot with a pin. Colored blob detection and averaging produces a

relatively accurate measurement of the 3D tip position in the camera frame. The robot is driven to 30 random locations within its workspace and its joint coordinates are recorded. When the four circles are visible in both color and depth images and agree on the tip position within 5 mm, we consider the tip to be accurately measured. A least-squares transformation matrix between joint coordinates and sensed positions is then fit to the data. After calibration we find the mean average error in the range of 0.8mm-1mm with the maximum error less than 3mm.

6.5.2 Measurement of system accuracy

We built an apparatus to measure the cumulative positioning error of the system including human set-up error, calibration error, localization error, and hardware inaccuracies. We built a button panel with a “virtual button” in the center, whose coordinates are at the center of two colored diagonal visual features (Fig. 6.7). We conducted 500 test pushes using panel localization to determine the button location. Between each push we changed the position and orientation of the panel, with the entire panel oriented on each of its four sides and with yaw altered to up to 45° . The tip position was measured using the calibration cross marker, and “true” button center position was measured using the larger cross features. Results show the average euclidean distance from the tip center to the sensed button center is 1.9mm, with an outlier of 8.39mm (Fig. 6.8). This outlier was likely caused by a hardware fault on the linear actuator.

To isolate performance of our panel localization system we performed 500 localization readings on a different test panel that has several buttons (slider, rocker, push button, and switch) and a color-based fiducial on the push button. We compare the estimated button position from panel localization against the “ground truth” location from color tracking. Testing repeatability for a static panel (i.e., the effect of camera noise) gives a maximum euclidean distance error of 2.56mm using only QR localization, which is reduced to 0.5mm with ICP. Repeating this experiment while shifting and rotating the panel between each reading, localization with QR code only had a maximum error 2.76mm, reduced to 1.3mm



Figure 6.7: The above apparatus is used for testing the accuracy of the system: (a) Robot and camera are calibrated using a cross marker, (b) Center button is located at the intersection of opposing diagonal circles, (c) Center plate is removed during the test to not interfere with tip motion.

with ICP.

6.5.3 Test panel experiments

As a controlled test of our platform’s reliability in operating switches, we built 2 test panels, one with toggle switches and another with push buttons. The push button panel consists of 5mm radius buttons made with polished plastic. Each pair of buttons is separated by at least 2mm. The toggle test panel holds 5 toggles of different size, shape, and operating force ranging from 2.5N to 8.2N. All buttons are located within a 12cm by 10cm rectangle centered on a 30cm by 23cm flat panel, placed approximately 21cm away from the robot’s base.

In the switch test, the robot localizes and switches on the 5 switches in sequence and then reports the perceived state of the switches. Then, it switches them off and again reports their perceived state. In case of a failure, we manually flip the switch to the desired state for the next sequence. For the pushbutton test, robot localizes and pushes 10 buttons in sequence. In each test, the robot runs 10 sequences and attempts 100 actuations in total. The panel is illuminated with indoor office lighting. To judge the impact of different

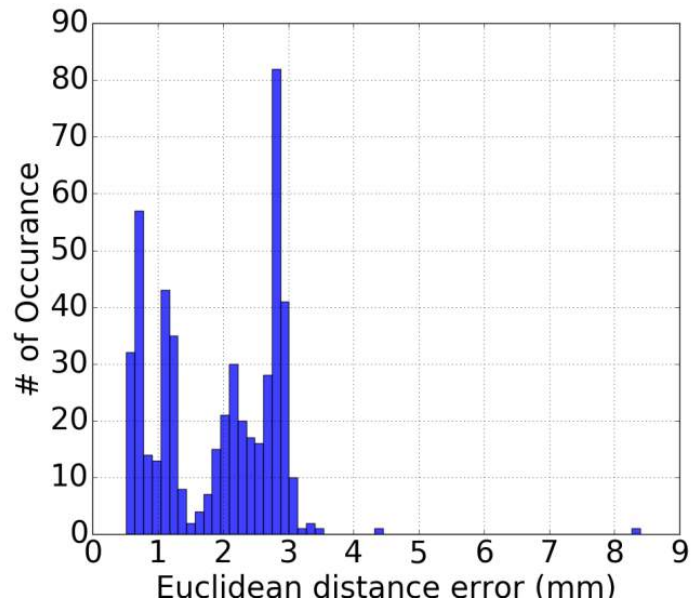


Figure 6.8: Histogram of system errors.




Figure 6.9: Error distribution in XY.

Table 6.2: Test Panel Experiments

Button Type	Toggle	Toggle	Toggle	Toggle	Push	Push	Push	Push
Tip	Cylinder	Scallop	Cylinder	Scallop	Cylinder	Scallop	Cylinder	Scallop
ICP	Off	Off	On	On	Off	Off	On	On
Success rate	79%	96%	84%	99%	98%	99%	98%	100%
State detection	100%	100%	100%	100%	n/a	n/a	n/a	n/a
Duration (min)	20	20	22	22	13	13	17	16
Time / push (s)	12	12	13.2	13.2	7.8	7.8	10.2	9.6

components on performance, we performed these tests with the two tips (cylindrical and scallop) and with and without ICP activated.

Tab. 6.2 shows that although the cylindrical tip performed well at button pushing, it failed in roughly one fifth of the switch attempts. We observed that errors occurred due to slippage or flexing of the robot structure. The scallop tip eliminated slippage, although it still failed 4 times when activating the lower left button, which had a relatively high 8.2N OF. The failure case was further reduced to only 1 when ICP was used in localization. The scallop tip design still has a high success rate on push buttons, even though it has a larger cross section than the small buttons we tested on. This is due to the curved tip that can direct force within a small area.

6.5.4 Experiment pressing in an office building and home

We tested our system on all accessible buttons on doors/walls and large electrical appliances in a 4-story office building and a 2 story residential town house. We mounted our robot on an adjustable-height stand with lockable wheels and pushed the robot along the corridor while testing on all buttons that are accessible and safe to test on.



Figure 6.10: A selection of button panels that SwitchIt succeeded in activating: (a) Door passcode, (b) Disabled door exit, (c) Thermostat buttons, (d) Switch with dimmer, (e) Office stereo control, (f) Security door entrance, (g) Electronic keypad, (h) Light switch

We have tested on 98 different button panels that contain a total of 379 individual buttons. The test set covers 39 distinct classes of panels and all 6 button types. We tested all distinct button on every panel, but when panels contain many identical buttons, we did not test every button. Specifically, if a panel contains exact duplicates of one button (such as a numeric keypad), we only tested 2 or 3 of them at extreme positions. In total, we asked the robot to operate 234 buttons.

Our robot succeeded in activating 224 of 234 buttons at the first attempt. We note that several of these buttons were quite challenging (Fig. 6.10). Successes include office passcode entry with small, stiff, and slippery metallic buttons; small and stiff rockers; turn knobs on a classroom stereo control; and non-conventional light switches that are activated from the side.



Figure 6.11: Four button types that SwitchIt failed to activate: (a) An old-style timer, (2) An elevator pull button, (3) A washer control button, (4) A turn knob on oven.

6.5.5 Test panel experiments

Out of the 10 failures, 3 switches are within the capability of the device, but failed due to various positioning errors. The other 7 were of 4 button designs our system currently cannot actuate (Fig. 6.11). They include “push to stop, pull to run” emergency stop button with a shallow smooth indentation for human finger to pinch. The hook in our finger cannot establish a solid hold onto this indentation. The other 3 are variations of turn knobs such as an old-style timer that requires more than 20N to turn, 3 oven temperature knobs that must be pushed in while turning, and 2 washing machine controls that are “pull to start, push to stop, and turn to select.” More work is needed to develop actuation strategies for these atypical button types.

Chapter 7

Future Work

An automated robot warehouse brings possible benefits such as increased up-time, higher total throughput, and lower accident rates. In particular, autonomous dense packing of arbitrary objects improves the storage capacity, decreases the delivery cost and saves packing materials. This thesis proposes a formulation of the packing problem that is tailored to the automated warehousing domain. Besides minimizing waste space inside a container, the problem requires stability of the object pile during packing and the feasibility of the robot motion executing the placement plans. To address this problem, a set of stability and robot packability constraints are formulated, and a constructive packing pipeline is proposed to pack geometrically complex, non-convex objects while satisfying these constraints. To evaluate the proposed planner under real-world uncertainties such as vision, grasping and modeling errors, a systematic evaluation on a state-of-the-art physical packing testbed is performed to study the sources of error and models their magnitude. Exhaustive experiments are conducted in Monte Carlo simulation and on the physical testbed to examine the feasibility of the packing placements, under open-loop baseline condition as well as two strategies for improving the robustness of robotic packing. Empirical results demonstrates the proposed planner produces stable and high-quality packing plans compared with other 3D packing methods, and success rate of up to 98% can be achieved on a physical robot when using robustness measures despite cascading real-world uncertainties.

In Chapter 2, a constructive pipeline is developed that can pack geometrically complex, non-convex objects with stability while satisfying robot constraints. A new Heightmap-Minimization heuristic is proposed as a positioning heuristic for efficient 3D irregular shape packing. Simulation results on exhaustive datasets demonstrate the effectiveness of the pipeline and the advantage of the new heuristic in finding stable and robot-packable plans. Robot-packable plans are shown to be far more successful in open-loop execution than

non-overlap methods used in prior work.

In chapter 3, two novel packing problems with nondeterministic item ordering are formulated. We presented practical solvers that handle irregular 3D shapes and item sets up to size 10. This work opens up several interesting theoretical and practical questions, such as the minimal number of plans needed to guarantee NDOP coverage, whether complete QOP algorithms exist for rectilinear items, whether efficiency gains are possible with multiple identical items, and whether restrictions on item shape can overcome exponential worst-case complexity.

In chapter 4, we started implementation the offline algorithm onto a physical system. Since the proposed algorithm assumes known 3D models of objects for planning and collision detection, and a highly precise model lead to better planning. We developed an in-hand object scanning pipeline that is easy to use and obtain results swiftly. Experiments demonstrate that the novel BackFlow video segmentation technique leads to improved tracking accuracy for fast changing objects when compared to other state-of-the-art methods, which leads to better reconstruction of 3D models. The only equipment needed by the method is an RGB-D camera and standard PC. Moreover, the software is easy to use for novices, taking approximately 2 minutes of manual input per object.

In chapter 5, we analyzed multiple sources of error in an integrated, automatic robotic packing system, and presented two strategies for overcoming these errors to increase packing success rates. In the robust planning strategy, geometric collision margins were added to avoid inadvertent contact. In the closed-loop strategy, vision was used to sense the state of the object pile, and to push items poking out of the container, and to replan if items slipped from their planned positions. The combination of both strategies led to higher success rates than either alone, and much higher success than open-loop planning. In all strategies, force feedback was also important to ensure items were successfully grasped and to avoid crushing items.

Future work could address non-rigid objects or uncertainty in 3D scanned models. We may also be able to increase the execution success rate by implementing a more conservative

stability check, or to perform a more sophisticated closed-loop execution. One method may be to check multiple perturbed gravity directions for feasibility [OR06], but it is still a challenge to ensure robustness with varying centers of mass and surface friction. We are also interested in improving replanning stability tests by identifying the shifted configuration of the object pile rather than treating it as a fixed, infinite-mass object. Finally, we are very interested in the use of manipulation to re-position items after they have been placed, such as compressing items toward a corner to make more space, or to repack items that have fallen.

Bibliography

- [AAB18] Steven Diamond Akshay Agrawal, Robin Verschueren and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- [AAD07] Pedram Azad, Tamim Asfour, and R. Dillmann. Stereo-based 6d object localization for grasping with humanoid robot systems. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 919–924, 2007.
- [AB94] R. Adams and L. Bischof. Seeded region growing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(6):641–647, June 1994.
- [ARC] Amazon robotics challenge official rules.
- [Art66] Jr Art, Richard Carl. *An approach to the two dimensional irregular cutting stock problem*. PhD thesis, Massachusetts Institute of Technology, 1966.
- [ASS⁺12] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012.
- [BCR80] B. Baker, E. Coffman, Jr., and R. Rivest. Orthogonal packings in two dimensions. *SIAM Journal on Computing*, 9(4):846–855, 1980.
- [BDD95] H. Bruyninckx, S. Dutre, and J. De Schutter. Peg-on-hole: a model based solution to peg and hole alignment. In *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, volume 2, pages 1919–1924 vol.2, 1995.
- [BETG08] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer Vision and Image Understanding*, 110(3):346 – 359, 2008. Similarity Matching in Computer Vision and Multimedia.
- [Bha] Rajen Bhatt. Skin segmentation data set. “<https://archive.ics.uci.edu/ml/datasets/skin+segmentation>. Accessed: 2018-09-12.
- [BHKW10] E. K. Burke, R. S. R. Hellier, G. Kendall, and G. Whitwell. Irregular packing using the line and arc no-fit polygon. *Operations Research*, 58(4-part-1):948–970, 2010.

- [BK04] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision, 2004.
- [BKP11] Abdeslam Boularias, Oliver Kroemer, and Jan Peters. Learning robot grasping from 3-d images with markov random fields. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1548–1553. IEEE, 2011.
- [BLP13] Julia A. Bennell, Lai Soon Lee, and Chris N. Potts. A genetic algorithm for two-dimensional bin packing with due dates. *International Journal of Production Economics*, 145(2):547 – 560, 2013.
- [BM92] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- [BMAK14] J. Bohg, A. Morales, T. Asfour, and D. Kragic. Data-driven grasp synthesis—a survey. *IEEE Transactions on Robotics*, 30(2):289–309, April 2014.
- [BR02] Fausto Bernardini and Holly Rushmeier. The 3d model acquisition pipeline. *Computer Graphics Forum*, 21(2):149–172, 2002.
- [C&17] C&K Switches. Product catalog. “<http://www.ckswitches.com/media/1416/ps.pdf>, 2017. (accessed March 2, 2017).
- [CB19] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. “<http://pybullet.org>, 2016–2019.
- [CBB⁺18] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman. Analysis and observations from the first amazon picking challenge. *IEEE Transactions on Automation Science and Engineering*, 15(1):172–188, Jan 2018.
- [CCC⁺08] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. *MeshLab: an Open-Source Mesh Processing Tool.*, volume 1. 01 2008.
- [CCSS01] Yung-Yu Chuang, B. Curless, D. H. Salesin, and R. Szeliski. A bayesian approach to digital matting. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 2, pages II–264–II–271 vol.2, 2001.

- [CFI⁺04] M. Callieri, A. Fasano, G. Impoco, P. Cignoni, R. Scopigno, G. Parrini, and G. Biagini. Roboscan: an automatic system for accurate and unattended 3d scanning. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004.*, pages 805–812, 2004.
- [CFL13] Camille Couprie, Clément Farabet, and Yann LeCun. Causal graph-based video segmentation. *CoRR*, abs/1301.1671, 2013.
- [CJ01] János Csirik and David S Johnson. Bounded space on-line bin packing: Best is better than first. *Algorithmica*, 31(2):115–138, 2001.
- [CJCH12a] S. Chitta, E. G. Jones, M. Ciocarlie, and K. Hsiao. Mobile manipulation in unstructured environments: Perception, planning, and execution. *IEEE Robotics Automation Magazine*, 19(2):58–71, June 2012.
- [CJCH12b] S. Chitta, E. G. Jones, M. Ciocarlie, and K. Hsiao. Mobile manipulation in unstructured environments: Perception, planning, and execution. *IEEE Robotics Automation Magazine*, 19(2):58–71, June 2012.
- [CLT⁺17] J. Cheng, S. Liu, Y.-H. Tsai, W.-C. Hung, S. Gupta, J. Gu, J. Kautz, S. Wang, and M.-H. Yang. Learning to segment instances in videos with spatial propagation network. *The 2017 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2017.
- [CM92] Yang Chen and Gérard Medioni. Object modelling by registration of multiple range images. *Image and Vision Computing*, 10(3):145 – 155, 1992.
- [CMPT⁺17] S. Caelles, K. K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. V. Gool. One-shot video object segmentation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5320–5329, 2017.
- [CPT08] Teodor Gabriel Crainic, Guido Perboli, and Roberto Tadei. Extreme point-based heuristics for three-dimensional bin packing. *INFORMS Journal on Computing*, 20(3):368–384, 2008.
- [CSB⁺17] Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, April 2017.
- [CTWY17] J. Cheng, Y. H. Tsai, S. Wang, and M. H. Yang. Segflow: Joint learning for video object segmentation and optical flow. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 686–695, 2017.

- [CW13] H. S. Chang and Y. C. F. Wang. Superpixel-based large displacement optical flow. In *2013 IEEE International Conference on Image Processing*, pages 3835–3839, 2013.
- [CWS⁺15] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set. *IEEE Robotics Automation Magazine*, 22(3):36–52, 2015.
- [DAV] DAVIS2017. Benchmark video object segmentation on davis. “http://davischallenge.org/soa_compare.html. Accessed: 2018-09-12.
- [DB16] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [dBKM⁺05] Edgar den Boef, Jan Korst, Silvano Martello, David Pisinger, and Daniele Vigo. Erratum to “the three-dimensional bin packing problem”: Robot-packable and orthogonal variants of packing problems. *Operations Research*, 53(4):735–736, 2005.
- [DR19] Siyuan Dong and Alberto Rodríguez. Tactile-based insertion for dense box-packing. *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7953–7960, 2019.
- [Dra05] Samuel Drake. Using compliance in lieu of sensory feedback for automatic assembly. 08 2005.
- [DVD02] Kathryn A. Dowsland, Subodh Vaid, and William B. Dowsland. An algorithm for polygon placement using a bottom-left strategy. *European Journal of Operational Research*, 141(2):371 – 381, 2002.
- [Ege09] Jens Egeblad. Placement of two- and three-dimensional irregular shapes for inertia moment and balance. *International Transactions in Operational Research*, 16:789 – 807, 06 2009.
- [EK09] Leah Epstein and Elena Kleiman. Resource augmented semi-online bounded space bin packing. *Discrete Applied Mathematics*, 157(13):2785–2798, 2009.
- [ENO07] Jens Egeblad, Benny K. Nielsen, and Allan Odgaard. Fast neighborhood search for two- and three-dimensional nesting problems. *European Journal of Operational Research*, 183(3):1249 – 1266, 2007.
- [EW15] Stefan Edelkamp and Paul Wichern. Packing irregular-shaped objects for 3d printing. *Advances in Artificial Intelligence. Lecture Notes in Computer Science, vol 9324*, pages 45–58, 09 2015.

- [Far03] Gunnar Farneback. Two-frame motion estimation based on polynomial expansion. In Josef Bigun and Tomas Gustavsson, editors, *Image Analysis*, pages 363–370, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [fcl] Flexible-collision-library.
- [FPZ03] Oluf Faroe, David Pisinger, and Martin Zachariasen. Guided local search for the three-dimensional bin-packing problem. *INFORMS Journal on Computing*, 15(3):267–283, 2003.
- [Fre96] Jon W. Freeman. Hard random 3-sat problems and the davis-putnam procedure. *Artificial Intelligence*, 81(1):183 – 198, 1996.
- [GJ79] M. R. Garey and David S. Johnson. *Computers and intractability: A guide to the theory of np-completeness*. 1979.
- [GKHE10] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2141–2148, 2010.
- [GMM90] H. Gehring, K. Menschner, and M. Meyer. A computer-based heuristic for packing pooled shipment containers. *European Journal of Operational Research*, 44(2):277 – 288, 1990.
- [GMZ⁺99] Kenneth Y. Goldberg, Brian Mirtich, Yan Zhuang, John Craig, Brian Carlisle, and John F. Canny. Part pose statistics: estimators and experiments. *IEEE Trans. Robotics and Automation*, 15:849–857, 1999.
- [GO06] A. Miguel Gomes and José F. Oliveira. Solving irregular strip packing problems by hybridising simulated annealing and linear programming. *European Journal of Operational Research*, 171(3):811 – 829, 2006. Feature Cluster: Heuristic and Stochastic Methods in Optimization Feature Cluster: New Opportunities for Operations Research.
- [GPS89] D. M. Greig, B. T. Porteous, and A. H. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, 51(2):271–279, 1989.
- [GR80] J.A. George and D.F. Robinson. A heuristic for packing boxes into a container. *Computers and Operations Research*, 7(3):147 – 156, 1980.
- [HCT⁺11] Xin Han, Francis YL Chin, Hing-Fung Ting, Guochuan Zhang, and Yong Zhang. A new upper bound 2.5545 on 2d online bin packing. *ACM Transactions on Algorithms (TALG)*, 7(4):50, 2011.

- [HJJ⁺13] Hyeonjun Park, Ji-Hun Bae, Jae-Han Park, Moon-Hong Baeg, and Jae-heung Park. Intuitive peg-in-hole assembly strategy with a compliant manipulator. In *IEEE ISR 2013*, pages 1–5, 2013.
- [HPL17] Guoheng Huang, Chi-Man Pun, and Cong Lin. Unsupervised video co-segmentation based on superpixel co-saliency and region merging. *Multimedia Tools Appl.*, 76(10):12941–12964, May 2017.
- [HSKMG09] H. Hamer, K. Schindler, E. Koller-Meier, and L. V. Gool. Tracking a hand manipulating an object. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1475–1482, 2009.
- [HT01] E. Hopper and B. C. H. Turton. A review of the application of meta-heuristic algorithms to 2d strip packing problems. *Artificial Intelligence Review*, 16(4):257–300, Dec 2001.
- [HZRS15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- [IKH⁺11] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. Kinectfusion: Real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, pages 559–568, 2011.
- [JDU⁺74] D. Johnson, A. Demers, J. Ullman, M. Garey, and R. Graham. Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM Journal on Computing*, 3(4):299–325, 1974.
- [JPS13] Bonfim A. Junior, Plácido R. Pinheiro, and Rommel D. Saraiva. A hybrid methodology for nesting irregular shapes: Case study on a textile industry. *IFAC Proceedings Volumes*, 46(24):15 – 20, 2013. 6th IFAC Conference on Management and Control of Production and Logistics.
- [JR02] Michael J. Jones and James M. Rehg. Statistical color models with application to skin detection. *International Journal of Computer Vision*, 46(1):81–96, 2002.
- [kam88] Thomas kampke. Simulated annealing: Use of a new tool in bin packing. *Annals of Operations Research*, 16(1):327–332, Dec 1988.
- [KAO07] Jeong-Gwan Kang, Su-Yong An, and Se-Young Oh. Navigation strategy for the service robot in the elevator environment. In *Int. Conf. on Control, Automation and Systems*, pages 1092–1097. IEEE, 2007.

- [KCRN10] Ellen Klingbeil, Blake Carpenter, Olga Russakovsky, and Andrew Y Ng. Autonomous operation of novel elevators for robot navigation. In *IEEE Int. Conf. on Robotics and Automation*, pages 751–758. IEEE, 2010.
- [KGS⁺11] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. G2o: A general framework for graph optimization. In *2011 IEEE International Conference on Robotics and Automation*, pages 3607–3613, May 2011.
- [KH13] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 32(3):29:1–29:13, July 2013.
- [KHRF11] Michael Krainin, Peter Henry, Xiaofeng Ren, and Dieter Fox. Manipulator and object tracking for in-hand 3d object modeling. *Int. J. Rob. Res.*, 30(11):1311–1327, 2011.
- [kla] Klampt - intelligent motion laboratory at duke university.
- [KMV⁺16] M. Khodabandeh, S. Muralidharan, A. Vahdat, N. Mehrasa, E. M. Pereira, S. Satoh, and G. Mori. Unsupervised learning of supervoxel embeddings for video segmentation. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2392–2397, 2016.
- [KOM⁺03] Rie Katsuki, Jun Ota, Takahisa Mizuta, Tomomi Kito, Tamio Arai, Tsuyoshi Ueyama, and Tsuyoshi Nishiyama. Design of an artificial mark to determine 3d pose by monocular vision. In *IEEE Int. Conf. on Robotics and Automation*, volume 1, pages 995–1000. IEEE, 2003.
- [KOT⁺03] Rie Katsuki, Jun Ota, Yusuke Tamura, Takahisa Mizuta, Tomomi Kito, Tamio Arai, Tsuyoshi Ueyama, and Tsuyoshi Nishiyama. Handling of objects with marks by a robot. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, volume 1, pages 130–135. IEEE, 2003.
- [LBRF11] K. Lai, L. Bo, X. Ren, and D. Fox. A large-scale hierarchical multi-view rgb-d object dataset. In *2011 IEEE International Conference on Robotics and Automation*, pages 1817–1824, May 2011.
- [LISD10] H. Liao, T. Inomata, I. Sakuma, and T. Dohi. 3-d augmented reality for mri-guided surgery using integral videography autostereoscopic image overlay. *IEEE Transactions on Biomedical Engineering*, 57(6):1476–1486, 2010.
- [LKH⁺13] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *2013 IEEE International Conference on Computer Vision*, pages 2192–2199, 2013.

- [LLCY15] Xiao Liu, Jia-min Liu, An-xi Cao, and Zhuang-le Yao. Hape3d—a new constructive algorithm for the 3d irregular packing problem. *Frontiers of Information Technology & Electronic Engineering*, 16(5):380–390, May 2015.
- [LLTK14] Seok Ju Lee, Jongil Lim, Girma Tewolde, and Jaerock Kwon. Autonomous tour guide robot by using ultrasonic range sensors and qr code recognition in indoor environment. In *IEEE Int. Conf. on Electro/Information Technology (EIT)*, pages 410–415. IEEE, 2014.
- [LMO⁺16] Xiaoping Liao, Junyan Ma, Chengyi Ou, Fengying Long, and Xiangsha Liu. Visual nesting system for irregular cutting-stock problem based on rubber band packing algorithm. *Advances in Mechanical Engineering*, 8(6):1687814016652080, 2016.
- [LMV02] Andrea Lodi, Silvano Martello, and Daniele Vigo. Heuristic algorithms for the three-dimensional bin packing problem. *European Journal of Operational Research*, 141(2):410 – 420, 2002.
- [LMV04] Andrea Lodi, Silvano Martello, and Daniele Vigo. Tspack: A unified tabu search code for multi-dimensional bin packing problems. *Annals of Operations Research*, 131(1):203–213, Oct 2004.
- [Low04a] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [Low04b] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, pages 60–91, 2004.
- [LPC⁺00] Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Ginzton, Sean Anderson, James Davis, Jeremy Ginsberg, Jonathan Shade, and Duane Fulk. The digital michelangelo project: 3d scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 131–144, 2000.
- [LPK⁺18] Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4-5):421–436, 2018.
- [LPKQ16] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *CoRR*, abs/1603.02199, 2016.

- [LTHF17] Y. G. Lee, Z. Tang, J. N. Hwang, and Z. Fang. Inter-camera tracking based on fully unsupervised online learning. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 2607–2611, 2017.
- [MPV00] Silvano Martello, David Pisinger, and Daniele Vigo. The three-dimensional bin packing problem. *Operations Research*, 48(2):256–267, 2000.
- [MS15] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928, 2015.
- [MV98] Silvano Martello and Daniele Vigo. Exact solution of the two-dimensional finite bin packing problem. *Management Science*, 44(3):388–399, 1998.
- [MWG⁺10] W. Meeussen, M. Wise, S. Glaser, S. Chitta, C. McGann, P. Mihelich, E. Marder-Eppstein, M. Muja, V. Eruhimov, T. Foote, J. Hsu, R. B. Rusu, B. Marthi, G. Bradski, K. Konolige, B. Gerkey, and E. Berger. Autonomous door opening and plugging in with a personal robot. In *IEEE Int. Conf. on Robotics and Automation*, pages 729–736, 2010.
- [Naw28] B. Nawahi. Acoustic production in liquid filled ceramic environments. *J. Chem. Phys.*, 108:9893–9904, 1928.
- [NDRK09] Hai Nguyen, Travis Deyle, Matt S Reynolds, and Charles C Kemp. Pps-tags: Physical, perceptual and semantic tags for autonomous mobile manipulation. Georgia Institute of Technology, 2009.
- [Ngu88] Van-Duc Nguyen. Constructing force- closure grasps. *The International Journal of Robotics Research*, 7(3):3–16, 1988.
- [NSLV86] A.Y.C. Nee, K.W. Seow, S.L. Long, and V.C. Venkatesh. Designing algorithm for nesting irregular shapes with and without boundary constraints. *CIRP Annals*, 35(1):107 – 110, 1986.
- [OKA11] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints. In *2011 International Conference on Computer Vision*, pages 2088–2095, 2011.
- [OMR12] OMRON. Product catalog. “<http://www.ckswitches.com/media/1416/ps.pdf>, 2012. (accessed September 9, 2012).

- [OMR17] OMRON. Switches: Basic operating characteristics. “<https://www.omron.com/ecb/products/sw/special/switch/basic01-08.html>, 2017. (accessed March 2, 2017).
- [OR06] Yizhar Or and Elon Rimon. Computation and graphical characterization of robust multiple-contact postures in two-dimensional gravitational environments. *The International Journal of Robotics Research*, 25(11):1071–1086, 2006.
- [PASV15] J. Prankl, A. Aldoma, A. Svejda, and M. Vincze. Rgb-d object modelling for object recognition and tracking. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 96–103, 2015.
- [PGD⁺17] Deepak Pathak, Ross Girshick, Piotr Dollár, Trevor Darrell, and Bharath Hariharan. Learning features by watching objects move. In *CVPR*, 2017.
- [PGM⁺19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [PKA15] Paschalis Panteleris, Nikolaos Kyriazis, and Antonis A. Argyros. 3d tracking of human hands in interaction with unknown objects. In *BMVC*, 2015.
- [PKB⁺17] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3491–3500, 2017.
- [PLC⁺12] A. Prest, C. Leistner, J. Civera, C. Schmid, and V. Ferrari. Learning object class detectors from weakly annotated video. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3282–3289, 2012.
- [PLS13] A. Pratkanis, A. E. Leeper, and K. Salisbury. Replacing the office intern: An autonomous coffee run with a mobile manipulator. In *IEEE Int. Conf. on Robotics and Automation*, pages 1248–1253, 2013.

- [PPTM⁺16] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 724–732, 2016.
- [PS03] Soon-Yong Park and Murali Subbarao. An accurate and fast point-to-plane registration technique. *Pattern Recognition Letters*, 24(16):2967 – 2976, 2003.
- [PS14] S. Poullot and S. Satoh. Vabcut: A video extension of grabcut for unsupervised video foreground object segmentation. In *2014 International Conference on Computer Vision Theory and Applications (VISAPP)*, volume 2, pages 362–371, 2014.
- [PZK17] J. Park, Q. Zhou, and V. Koltun. Colored point cloud registration revisited. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 143–152, Oct 2017.
- [RC11] R. B. Rusu and S. Cousins. 3d is here: Point cloud library (pcl). In *2011 IEEE International Conference on Robotics and Automation*, pages 1–4, May 2011.
- [RD05] E. Rosten and T. Drummond. Fusing points and lines for high performance tracking. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 2, pages 1508–1515 Vol. 2, 2005.
- [RHHL02] Szymon Rusinkiewicz, Olaf Hall-Holt, and Marc Levoy. Real-time 3d model acquisition. *ACM Trans. Graph.*, 21(3):438–446, July 2002.
- [RI94] R.L. Rao and S.S. Iyengar. Bin-packing by simulated annealing. *Computers and Mathematics with Applications*, 27(5):71 – 82, 1994.
- [RKB04] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ”grabcut”: Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, August 2004.
- [RL01a] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, pages 145–152, 2001.
- [RL01b] Szymon Rusinkiewicz and Marc Levoy. Efficient variants of the icp algorithm. *Efficient Variants of the ICP Algorithm*, pages 145–152, 02 2001.
- [RMF12] Alberto Rodriguez, Matthew T Mason, and Steve Ferry. From caging to grasping. *The International Journal of Robotics Research*, 31(7):886–900, 2012.

- [RRKB11] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *2011 International Conference on Computer Vision*, pages 2564–2571, 2011.
- [RT00] M. A. Ruzon and C. Tomasi. Alpha estimation in natural images. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No.PR00662)*, volume 1, pages 18–25 vol.1, 2000.
- [RT16] H. Ramadan and H. Tairi. Automatic human segmentation in video using convex active contours. In *2016 13th International Conference on Computer Graphics, Imaging and Visualization (CGiV)*, pages 184–189, 2016.
- [Rut15] Rutgers APC RGB-D Dataset, 2015. (Accessed Jan 28, 2019).
- [San05] F. Sanford. Oxidized ferrous materials. *Journal of Junk*, 5(4):324–345, 2005.
- [Sei02] Steven S Seiden. On the online bin packing problem. *Journal of the ACM (JACM)*, 49(5):640–671, 2002.
- [SH70] M. Smith and I. A. Hall. *Handbook of Interstellar Travel*. Dover, New York, 7th edition, 1970.
- [SLMG⁺18] Max Schwarz, Christian Lenz, German Martin Garcia, Seong-Yong Koo, Arul Selvam Periyasamy, Michael Schreiber, and Sven Behnke. Fast object learning and dual-arm coordination for cluttered stowing, picking, and packing. 05 2018.
- [SMZ⁺16] Srinath Sridhar, Franziska Mueller, Michael Zollhoefer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d input. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2016.
- [SS10] Vladimir Sukhoy and Alexander Stoytchev. Learning to detect the functional components of doorbell buttons using active exploration and multimodal correlation. In *IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, pages 572–579. IEEE, 2010.
- [SSB⁺12] Guntram Scheithauer, Yu. G. Stoyan, Julia A. Bennell, Tatiana Romanova, and Aleksandr Pankratov. Containment of a pair of rotating objects within a container of minimal area or perimeter. 2012.
- [ST08] Ufuk Sakarya and Ziya Telatar. Graph-based multilevel temporal video segmentation. *Multimedia Systems*, 14(5):277–290, Nov 2008.

- [STD09] Giovanna Sansoni, Marco Trebeschi, and Franco Docchio. State-of-the-art and applications of 3d imaging sensors in industry, cultural heritage, medicine, and criminal investigation. *Sensors*, 9:568–601, 2009.
- [STS⁺19] Rahul Shome, Wei N. Tang, Changkyu Song, Chaitanya Mitash, Hristiyan Kourtev, Jingjin Yu, Abdeslam Boularias, and Kostas E. Bekris. Towards robust product packing with a minimalistic end-effector. *CoRR*, abs/1903.00984, 2019.
- [SUB96] Marcos Salganicoff, Lyle H Ungar, and Ruzena Bajcsy. Active learning for vision-based robot grasping. *Machine Learning*, 23(2-3):251–278, 1996.
- [TBS⁺16] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *Int. J. Comput. Vision*, 118(2):172–193, 2016.
- [TG17] Dimitrios Tzionas and Juergen Gall. 3d object reconstruction from hand-object interactions. *CoRR*, abs/1704.00529, 2017.
- [TK14] Sugitha. T.S. and Gladwin Jos K.T. Optical flow based on feature match and super pixel segmentation. In *2014 Fourth International Conference on Advances in Computing and Communications*, pages 243–246, Aug 2014.
- [TYB16] Y. H. Tsai, M. H. Yang, and M. J. Black. Video segmentation via object flow. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3899–3908, 2016.
- [VL17] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*, 2017.
- [VSA03] Vladimir Vezhnevets, Vassili Sazonov, and Alla Andreeva. A survey on pixel-based skin color detection techniques. In *IN PROC. GRAPHICON-2003*, pages 85–92, 2003.
- [VT03] Christos Voudouris and Edward P. K. Tsang. *Guided Local Search*, pages 185–218. Springer US, Boston, MA, 2003.
- [VVHS15] Joaquim L. Viegas, Susana M. Vieira, Elsa M. P. Henriques, and João M. C. Sousa. A tabu search algorithm for the 3d bin packing problem in the steel industry. In António Paulo Moreira, Aníbal Matos, and Germano Veiga, editors, *CONTROLO’2014 – Proceedings of the 11th Portuguese Conference on Automatic Control*, pages 355–364, Cham, 2015. Springer International Publishing.

- [WGC⁺10] Lei Wang, Songshan Guo, Shi Chen, Wenbin Zhu, and Andrew Lim. Two natural heuristics for 3d packing with practical loading constraints. In Byoung-Tak Zhang and Mehmet A. Orgun, editors, *PRICAI 2010: Trends in Artificial Intelligence*, pages 256–267, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [WH19a] F. Wang and K. Hauser. In-hand object scanning via rgb-d video segmentation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3296–3302, 2019.
- [WH19b] F. Wang and K. Hauser. Stable bin packing of non-convex 3d objects with a robot manipulator. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8698–8704, 2019.
- [WH19c] Fan Wang and Kris Hauser. Robot packing with known items and nondeterministic arrival order. In *Proceedings of Robotics: Science and Systems*, FreiburgimBreisgau, Germany, June 2019.
- [WHLC10] Wen-June Wang, Cheng-Hao Huang, I-Hsian Lai, and Han-Chun Chen. A robot arm for pushing elevator buttons. In *SICE Annual Conference 2010, Proceedings of*, pages 1844–1848. IEEE, 2010.
- [WJL⁺05] B. Wang, L. Jiang, J. W. Li, H. G. Cai, and H. Liu. Grasping unknown objects based on 3d model reconstruction. In *Proceedings, 2005 IEEE/ASME International Conference on Advanced Intelligent Mechatronics.*, pages 461–466, 2005.
- [WLG08] T. Weise, B. Leibe, and L. Van Gool. Accurate and robust registration for in-hand modeling. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [WSP15] Wenguan Wang, Jianbing Shen, and F. Porikli. Saliency-aware geodesic video object segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3395–3402, 2015.
- [WWLG09] T. Weise, T. Wismer, B. Leibe, and L. Van Gool. In-hand scanning with online loop closure. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1630–1637, 2009.
- [WWLG11] Thibaut Weise, Thomas Wismer, Bastian Leibe, and Luc Van Gool. Online loop closure for real-time interactive 3d scanning. *Computer Vision and Image Understanding*, 115(5):635 – 648, 2011.
- [WXZ⁺19] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densfusion: 6d object pose estimation

- by iterative dense fusion. In *Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [XL16] F. Xiao and Y. J. Lee. Track and segment: An iterative unsupervised approach for video object proposals. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 933–942, 2016.
- [XSNF17] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *CoRR*, abs/1711.00199, 2017.
- [YCB19] YCB Benchmarks - Object and Model Set, 2019. (Accessed Jan 28, 2019).
- [YFD⁺16] Kuan-Ting Yu, Nima Fazeli, Nikhil Chavan Daffe, Orion Taylor, Elliott Donlon, Guillermo Diaz Lankenau, and Alberto Rodriguez. A summary of team mit’s approach to the amazon picking challenge 2015. *CoRR*, abs/1604.03639, 2016.
- [YR18] Kuan-Ting Yu and Alberto Rodriguez. Realtime state estimation with tactile and visual sensing for inserting a suction-held object. *CoRR*, abs/1803.08014, 2018.
- [ZH04] Defu Zhang and Wenqi Huang. A simulated annealing algorithm for the circles packing problem. In Marian Bubak, Geert Dick van Albada, Peter M. A. Sloot, and Jack Dongarra, editors, *Computational Science - ICCS 2004*, pages 206–214, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [ZPK18] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018.
- [ZSQ99] B. D. Zarit, B. J. Super, and F. K. H. Quek. Comparison of five color models in skin pixel classification. In *Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems, 1999. Proceedings. International Workshop on*, pages 58–63, 1999.
- [ZSY⁺17] Andy Zeng, Shuran Song, Kuan-Ting Yu, Elliott Donlon, Francois Robert Hogan, Maria Bauzá, Daolin Ma, Orion Taylor, Melody Liu, Eudald Romo, Nima Fazeli, Ferran Alet, Nikhil Chavan Daffe, Rachel Holladay, Isabella Morona, Prem Qu Nair, Druck Green, Ian Taylor, Weber Liu, Thomas A. Funkhouser, and Alberto Rodriguez. Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. *CoRR*, abs/1710.01330, 2017.

- [ZSY⁺18] A. Zeng, S. Song, K. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, N. Fazeli, F. Alet, N. C. Daffe, R. Holladay, I. Morena, P. Qu Nair, D. Green, I. Taylor, W. Liu, T. Funkhouser, and A. Rodriguez. Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3750–3757, 2018.