Master's Thesis

Model-free Sensorless Manipulation

Pragna Mannam

CMU-RI-TR-19-06

May 2019

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Matthew T. Mason, Chair
Oliver Kroemer
Nancy Pollard
Adithya Murali

Submitted in partial fulfillment of the requirements
for the degree of Master of Science.

Keywords: manipulation; probabilistic reasoning, automation; manufacturing and logistics

To My Family

Abstract

This thesis is a study of 2D manipulation without sensing and planning, by exploring the effects of unplanned randomized action sequences on 2D object pose uncertainty. Our approach uses sensorless reorienting of an object to achieve a determined pose, regardless of the initial pose. Without using sensors and models of the object's properties, this work shows that under some circumstances, a long enough sequence of random actions will also converge toward a determined final pose of the object. This is verified through several simulation and real robot experiments where randomized action sequences are shown to reduce entropy of the object pose distribution. The effects of varying object shapes, action sequences, and surface friction properties are also explored.

Acknowledgments

I would like to first thank my advisor, Matt Mason, for giving me the opportunity to begin my journey in robotics and manipulation in my undergraduate years. This project has meant so much to me. I also want to thank Katharina Muelling, Oliver Kroemer, and Nancy Pollard for their mentorship and feedback on my thesis. Thank you to the entire manipulation lab for guiding me throughout the course of this project and for their fruitful discussion. I am grateful for my family, friends, and dissertation writing group for supporting me throughout this process. This work would not have been possible without them. Thank you.

Contents

List of Figures

Chapter 1

Introduction

Robots are envisioned to manipulate and interact with objects in unstructured environments and accomplish a diverse set of tasks. Towards this, reducing object pose uncertainty is often necessary for successful task execution. There are natural ways to reduce pose uncertainty including the addition of physical constraints, relative positioning to a known object's pose, and actively sensing the desired object's pose. In this paper, we explore a novel pose uncertainty reduction technique based on executing randomized sequence of actions. We evaluate our proposed pose uncertainty reduction technique on parts orienting, an industrial automation task.

Reducing task state uncertainty in parts orienting systems is an important part of factory automation, especially product assembly. The problem is to take parts in a disorganized jumble of multiple objects and to present them one at a time in a predictable pose. Most industrial solutions involve a part-specific mechanical design. One goal of parts-orienting research is to avoid part-specific mechanical designs, reducing the time required to develop the automation for a new or redesigned product.

Tray-tilting is one kind of part-agnostic object reorientation system. The original tray-tilting work was an early entry in a research approach termed "minimalism". Minimalism refers to "the art of doing X without Y", or "finding the minimal configuration of resources to solve a task" [5]. The purpose of the approach is not just to conserve resources, but to yield insights into the

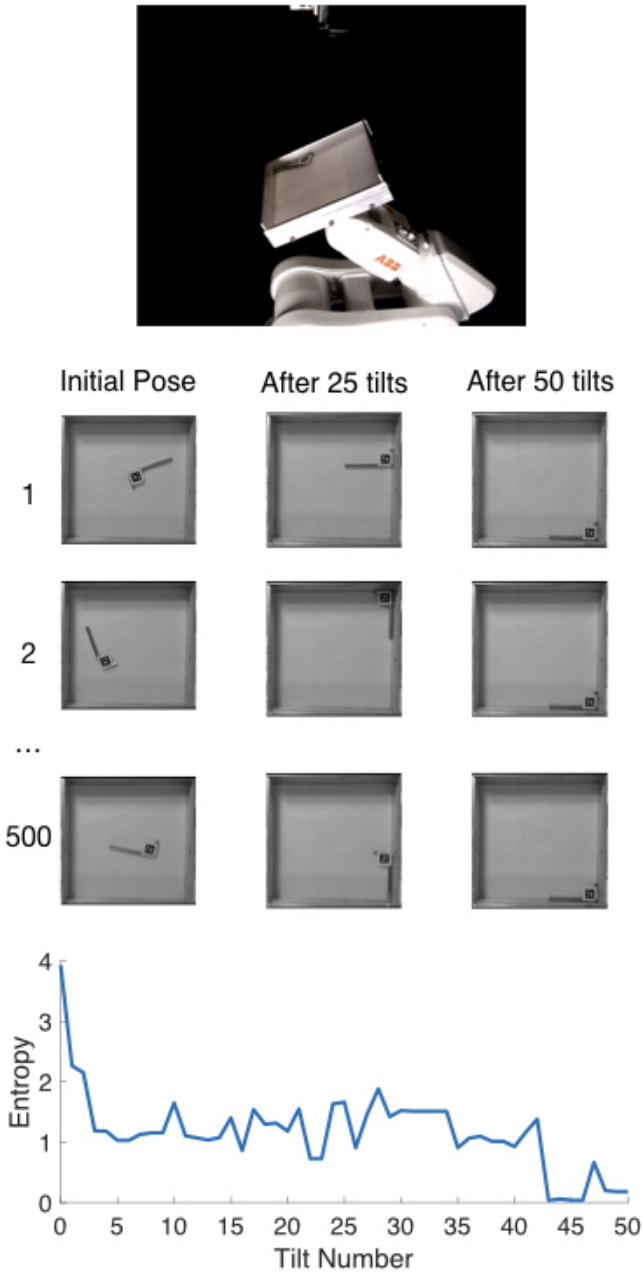

Figure 1.1: Experimental setup (top). An industrial robot tilts an allen key, with April Tag attached, in an aluminum tray. The overhead camera records the pose of the allen key after each tilt. Each trial (1), (2), . . . (500) performs the same random sequence of actions with a different initial position. The pose before the sequence, mid-sequence after 25 tilts, and after the sequence of 50 tilts are shown per trial, as well as the entropy of object pose distribution over 500 total repeated trials (bottom).

structure of tasks, the nature of perception, planning, and action, and robustness in simplicity.

The role of sensing in the sense-plan-act structure was examined by the tray-tilting work of Erdmann and Mason [10], which eliminated all uncertainty in the task state without sensing. Rather than sensing, the discrete set of feasible task states could sometimes be reduced to a singleton through a judicious choice of actions. Thus, in some tasks, task state uncertainty can be eliminated without sensing even with the noisy mechanics of frictional contact.

While the original paper by Erdmann and Mason [10] examined the role of sensing, this paper is an extension exploring similar minimalism in planning. We replace Erdmann and Mason's [10] planned sequence of actions with a randomized sequence of actions and evaluate the reduction in object pose uncertainty. Using tray-tilting random actions, instead of planning, can provide simple part-agnostic designs in factory automation. Our experiments stay close to the original work to focus on the role of planning. We test the limits of minimalism with respect to system complexity and hope to pursue its practical applications in future work. For this reason, we adopt the same task domain: planar sliding of a laminar object in a rectangular tray. The robot can tilt the tray as desired, and the goal is to move the object to a single final pose, irrespective of its initial pose. If independent actions do not scramble the task state by introducing too much disorder, then occasionally some action maps two initial task states to the same final task state. Furthermore, we expect the set of feasible task states to approach a singleton, for sufficiently long sequences, as seen in Figure 1.1. The phenomenon, while also reminiscent of contraction mapping, is similar to an interesting card trick called the Kruskal Count [1], so we have dubbed the phenomenon as the "Kruskal effect".

The goal of this paper is two-fold. The ultimate practical goal is online autonomous management of uncertainty in task state, in place of offline human engineering of the task domain. The immediate scientific goal is to better understand the stochastic nature of manipulation tasks. In particular, studying the evolution of entropy under a random sequence of actions is a previously unexplored approach that reveals something of the intrinsic nature of the task. Our results

have immediate implications for machine learning approaches, specifically for the problem of searching for an effective plan, or estimating the stochastic behavior of a given plan. To explore the Kruskal effect, we conduct an experiment under ideal conditions to explore what behavior we expect. Then, we discuss the results from simulation experiments using different random sequences, various triangular object shapes, and varied friction floor noise levels to test the tolerance of our system. Finally, we show the experimental results from running a random sequence of actions on a robotic arm. For objects similar to allen keys, relatively low tray friction noise, and a long enough sequence of random actions, we show that the Kruskal effect applies. We also observe that it does not apply as well to cases with high tray friction noise, and exploration into more cases is left for future work. The insights we gain from our exploration of the limits of Kruskal effect can lay the foundations for compartmentalized tray-tilting of a kit of parts in factory automation or 3D pose determination in future work.

1.1 Thesis Outline

First, Chapter 3 will discuss how we reduce object pose uncertainty and calculate the amount of pose uncertainty after every action in a sequence using entropy. A proof of concept example in Chapter 4 illustrates what we expect from randomized action sequences with respect to pose uncertainty. Experimental setup and results are presented in Chapter 5. Finally, we discuss our observations in Chapter 6 and conclude with directions for future research in Chapters 7 and 8.

Chapter 2

Previous Work

The problem of presenting a single object from disorganization has interested robotics researchers as far back as Grossman and Blasgen's work in 1975 [12]. Grossman and Blasgen introduced a fixed tilted tray that used vibration to eliminate the effects of friction. An irregular part in the tray would settle into one of a small number of stable poses, and the robot used a touch probe to disambiguate the pose. Várkonyi [19] includes additional details on approaches to the problem by using simulation to systematically evaluate various pose estimators.

Erdmann and Mason [10] substituted a fixed tray with an active tilting tray, and showed that for some parts, a sequence of tilts would reduce the possible poses to a singleton, completely orienting the part without a touch probe or any other sensor.

While the tray is not part-specific, the Erdmann and Mason [10] approach uses part-specific motions. In this paper, we substitute the motions with a random sequence of tilts, which is not part-specific. If we can identify an interesting class of parts that are oriented by a random sequence, then we have what is sometimes termed a "universal" parts orienting system. Böhringer et al. [5] includes an overview of universal parts orienting research, detailing the design and implementation of planar force vector fields that will orient asymmetric laminar parts.

Sanderson [18] introduced parts entropy in the context of automated manufacturing. We use probability density functions in the configuration space, $SE(2)$ for planar motion of rigid parts,

and we use entropy to measure and compare distributions. Our calculation of entropy is based on Chirikjian's work on computing the discrete entropy of histograms [7].

Previously, pose uncertainty has been addressed using mechanical constraints, rather than sensing, in manipulation. Brost [6] uses squeeze-grasp actions to intrinsically reduce uncertainty of the object's position. Goldberg [11] planned sequences of pushes and squeezes to orient planar polygons up to symmetry. Zhou et al. [20] plans similar sequences based on an efficient simulation of planar pushing. Berretty et al. [3] proposed an approach of executing pulling actions using overhead fingers for object reorientation. Akella and Mason [2] applied a similar approach to parts with uncertain shape. Unlike these previous works, we use a random sequence of actions to reduce the uncertainty associated with the pose of an object. A random sequence of actions is a part-agnostic plan that minimizes software complexity and hardware changes for new parts.

Furthermore, we can exploit mechanical constraints in the environment. Our randomized sequence of actions uses mechanical constraints to manipulate the task state. Similarly, external contacts can be exploited to reach a final task state using prehensile pushing and in-hand object reorientation [9, 16].

Chapter 3

Reducing Pose Uncertainty

Before, during, and after a manipulation task, it is important to have an estimate of where the object to manipulate is located. When a manipulation attempt fails, or the object is not where it was supposed to be, we can employ certain strategies to decrease the pose uncertainty.

Some of these strategies were detailed in Chapter 2, including vibration, squeeze-grasps, pushing, and pulling. One way to measure uncertainty in object pose is to compute entropy using the probability distribution of objects' poses at a given step over repetitions of the task. For example, we can throw a key on a table one hundred times and check the distribution of the key poses over the repeated trials. The amount of order or disorder in the object poses (position and orientation) shows us the object pose uncertainty after an action, such as throwing.

As an example, let's look at pushing actions on a small disk. In Figure 3.1, there are two possible actions, pushing two walls in horizontally or vertically. The small disk, starting in any random initial position, moves respectively with the pushing force. Since the set of actions only consists of two actions, any random sequence of non-redundant actions would reorient the small disk to the same final (non-predetermined) pose in Figure 3.1.

In our experiments, we put an object in a tray and tilt the tray in random directions to reorient the object to a final determined pose with no prior knowledge of the initial object pose or object model. The random actions are used to decrease object pose uncertainty.

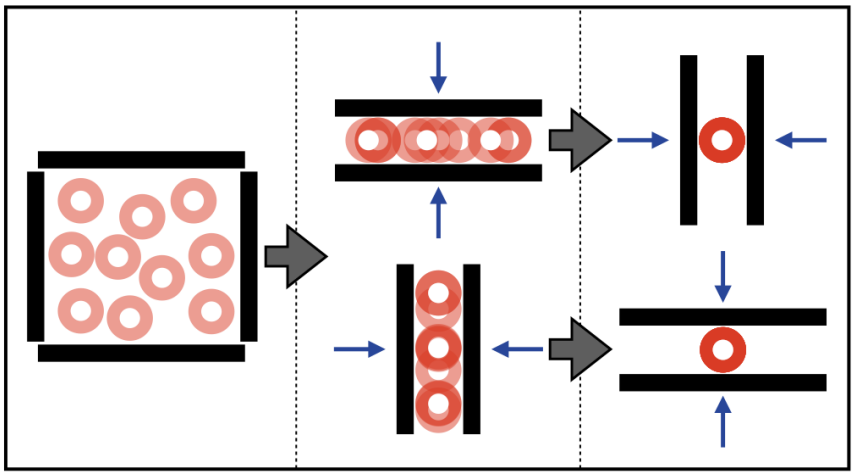

Figure 3.1: The effect of orthogonal actions on object pose. Long blocks represent a manipulator's two fingers with which we can execute horizontal and vertical squeeze grasps. Translucent disks indicate possible initial poses, and opaque disks indicate the resulting final poses after executing the action.

Measuring Entropy

Parts entropy describes the probability distribution of an object's pose over repeated tasks [18]. We measure object pose uncertainty using parts entropy throughout our randomized action sequences over many trials. Using parts entropy from Sanderson [18] and notation from Lee et al. [14], we define an object's pose in a tray of size $a \times b$ with the tuple (x, y, θ) where each coordinate is discretized with uniform spacing such that

$$x \in \{x_j : j = 1, \ldots, \alpha\} \quad \text{on} \quad [0, a] \tag{3.1}$$

$$y \in \{y_k : k = 1, \ldots, \beta\} \quad \text{on} \quad [0, b] \tag{3.2}$$

$$\theta \in \{\theta_m : m = 1, \ldots, \gamma\} \quad \text{on} \quad [0, 2\pi] \tag{3.3}$$

The number of discretized intervals are

$$\alpha = \frac{a}{\epsilon_p}, \quad \beta = \frac{b}{\epsilon_p}, \quad \gamma = \frac{2\pi}{\epsilon_r} \tag{3.4}$$

where ϵ_p and ϵ_r are the positional and rotational resolutions, respectively. We selected resolutions ϵ_p and ϵ_r such that α, β, and γ are integers. Object poses can only change through a set of tray

tilting actions A. Tilting directions were chosen to make a sequence S composed of N actions.

$$S = \{a_1, a_2, ..., a_N\}, a_i \in A$$

where A is the set of tilting actions in the cardinal directions. We execute a sequence S consisting of N random samples from A with replacement, and track the resulting sequence of object poses. We repeat the same sequence M times to obtain an estimated pose probability distribution after each action a_i,

$$f^i(x, y, \theta) = \frac{1}{M} V^i_{x,y,\theta} \tag{3.5}$$

where $V^i_{x,y,\theta}$ is the number of object poses that occupy the 3-dimensional voxel, (x, y, θ) after executing action a_i.

Given the pose probability distributions, we can compute the system entropy H^i following action a_i.

$$H^i = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \sum_{\theta \in \Theta} f^i(x, y, \theta) \log_2 f^i(x, y, \theta) \tag{3.6}$$

H can be interpreted as the number of additional information bits required to specify the object pose. If the pose distribution is uniform prior to the first tilt, then the entropy would be close to the logarithm of the number of voxels, $H^0 = \log_2(\alpha \times \beta \times \gamma)$. Ideally, the sequence converges to a fully determined pose, and the entropy drops to zero, $H^N = 0$. In terms of object pose uncertainty, high entropy corresponds to more uncertainty while low entropy corresponds to low uncertainty.

The main experimental challenge is the number of experiments required to reliably estimate the probability distribution of the poses. Lane [13] suggests that the number of trials M should satisfy

$$\alpha \times \beta \times \gamma = 2M^{1/3} \tag{3.7}$$

where $\alpha \times \beta \times \gamma$ is the number of voxels. The implication is that a large number of trials is required for even a very modest number of voxels. For our physical experiments (Section 5.3)

we selected a $3 \times 3 \times 3$ grid, which requires $M = 2460$ trials for a high-quality estimate of the probability distribution of the object poses. We used an action sequence consisting of $N = 50$ tilts, resulting in a total of $M \times N = 123,000$ tilts.

Unfortunately, the object and tray wear down after hundreds of tilts, changing the frictional properties of the system over time. We therefore settle with $M = 500$ trials and a total of $25,000$ tilts. As discussed in Chapter 6 and shown in Figure 6.1, the number of occupied voxels significantly decreases after the first few actions. After the first tilt, we do not occupy more than 10 voxels. According to Equation 3.7, we would need 125 trials for a maximum of 10 voxels to have reliable entropy calculations for all subsequent tilts. In effect, we have a much smaller number of occupied voxels, which leads us to believe that the smaller number of trials are sufficient for our experiments. To conduct experiments on a large scale without the real world challenges such as wear and tear, we look towards simulation.

For analyzing simulation data, we selected a $4 \times 4 \times 4$ grid, which would require $M \approx 32,000$ trials for high quality estimates of the probability distribution of the poses, according to Equation 3.7. Our three simulation experiments in Sections 5.2.1, 5.2.2, and 5.2.3 tested a total of 78 sequences ($M = 10,000$ trials per entropy trend for the first two experiments and $M = 1,000$ trials for the third). This results in almost $78 \times 32,000 \approx 2,500,000$ trials in total if we were to occupy all $4 \times 4 \times 4 = 64$ voxels across the tested action sequences. Instead, we conducted $600,000$ trials with sequences consisting of 50 actions resulting in $30,000,000$ tilts in simulation data.

Please note that the effect of our choice of voxel size is reduced by focusing on the change in entropy, rather than the absolute entropy [7].

Chapter 4

Proof of Concept

For proof of concept, let's discuss the reduction in object pose uncertainty to a single pose under ideal conditions. We explore this idea with a proof of concept experiment. Naturally, real experiments will change the results with the effects of Coulomb friction, gravity, material interactions, and wear and tear.

The experiment reorients an L-shaped object in a tray through tilting the tray in random directions (up, down, left, right, and diagonals). We empirically define the transition states assuming deterministic actions and stable poses after each random action. This was achieved through manual tray-tilting and observations of the resulting object poses. There are only 40

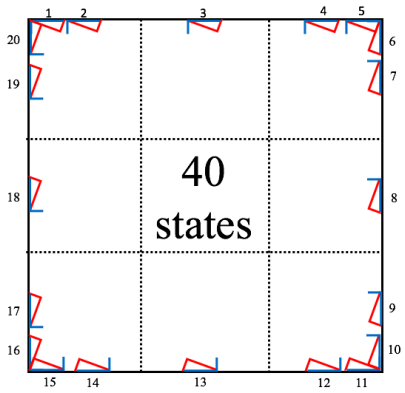

Figure 4.1: Possible states given ideal conditions of tray-tiling L-shaped object.

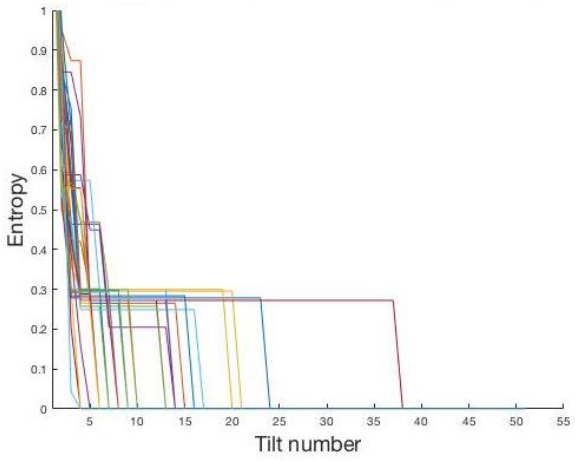

Figure 4.2: Entropy trend of object pose uncertainty over 1000 trials with ideal conditions. Randomized action sequences were 50 actions long and each line represents a newly randomized action sequence repeated 1000 times with different initial object orientations.

possible states with the object stably aligned to one or two walls as seen in Figure 4.1. The two stable orientations are illustrated in red and blue and there are twenty stable poses (aligned to a wall). Initialization requires the object to begin in one of these states and transition to any of the forty states as dictated by the transition matrix given the current state and action. We also assume that the effect of all actions is deterministic. This ensures that once an action causes two poses to converge to one pose, the object poses will never diverge throughout the rest of the action sequence.

After choosing a randomized sequence of actions, we can compute one entropy trend (line) of how the amount of object pose uncertainty changes over repetitions of that same sequence with different initial object poses. In Figure 4.2, many entropy trends are shown, each with a distinct random sequence. The entropy is monotonically decreasing because of the deterministic effects of our actions in this ideal experiment. We show that all random sequences reduce pose uncertainty, but at varying speeds. Thus, there is no way to determine if a sequence is long enough for zero object pose uncertainty because each random sequence will have different convergence rates. We use this proof of concept to show that we expect a decrease in object pose uncertainty using randomized sequences. Chapter 5 uses real-world dynamics (or approximations of them)

to show how we can use mechanical constraints to reduce object pose uncertainty in real world situations.

Chapter 5

Experiments

5.1 Experiment Setup

In order to measure the effects of randomized action sequences on objects, we start with a similar setup to Mason and Erdmann's sensorless manipulation experiment [10]. We tilt a tray in random directions instead of the original planned actions to reorient an allen key (also referred to as allen wrench or hex key). An allen key is used to test model-free sensorless manipulation since we know it has worked in planned object reorientation experiments [10]. The tray is attached to a commercial robot arm, and for our purposes we are using ABB IRB 120. It is important to note that a 6-DOF arm is not necessary to apply random actions to the tray-tilting setup as picking a direction of steepest descent only requires two degrees of freedom. An overhead camera records the pose of the allen key as it moves (shown in Figure 5.6). In simulation, we model the allen key as a two rectangular prisms connected by a sphere to make an 'L' shape with the same dimensions as the physical allen key. The set of actions are defined as

$$A = \{\texttt{N, NE, E, SE, S, SW, W, NW }\}$$

such that each action specifies the titling direction. If action '`N`' is executed, the tray's direction of steepest descent is north. Hence, the azimuth or direction of steepest ascent is south. A

vigorous vibration is used to randomly initialize the object pose before each sequence of actions. We have ensured that the distribution of initial poses covers the occupiable voxels uniformly. All sequences S are of length $N = 50$ actions, where there are only 8 possible actions. Thus, we have a total of $8^{50} \approx 1.43 \times 10^{45}$ possible sequences.

We want to ensure that there is not *just* one random sequence that reorients an object. Additionally, we want to explore what sort of friction levels and object shapes the system can handle until we can no longer expect random actions to result in object reorientation.

5.2 Simulation Experiments

Executing the experiment first in simulation enables us to generate the necessary number of trials required to estimate the object pose distribution with a sufficient pose resolution, across varied randomized action sequences, object shapes, and friction noise levels. For a realistic simulation we used the multibody contact friction model library in MATLAB Simscape. The tray used in the tray-tilting experiments was modeled as a box with no lid. The actions were 30-degree tray tilts in any of the eight cardinal directions. We started with an L-shaped object, mimicking the allen key used by Erdmann and Mason [10]. In the rest of the paper, we will use the terms actions and tilts interchangeably.

We simulate the contact model as a linear spring damped normal force with parameters selected to match experimentally observed metal-on-metal interactions. In Section 5.2.2, we used a few other polygonal shapes, as shown in Figure 5.1. All object interactions were modeled similarly, even for varied object shapes. The friction model is stick-slip with a velocity threshold [15]. To simulate noise during sliding for the tray friction noise in Section 5.2.3, the coeffecient of friction is varied spatially with an amplitude that we can vary to explore the effect of different friction noise levels.

The initial object pose in each trial was sampled from a uniform distribution in the object's configuration space (CSpace). Samples where the object was in collision with the wall were

rejected. To compute entropy throughout the sequence of actions as described in Section 3, we discretized the CSpace (x, y, θ) into $4 \times 4 \times 4 = 64$ (x, y, θ) voxels.

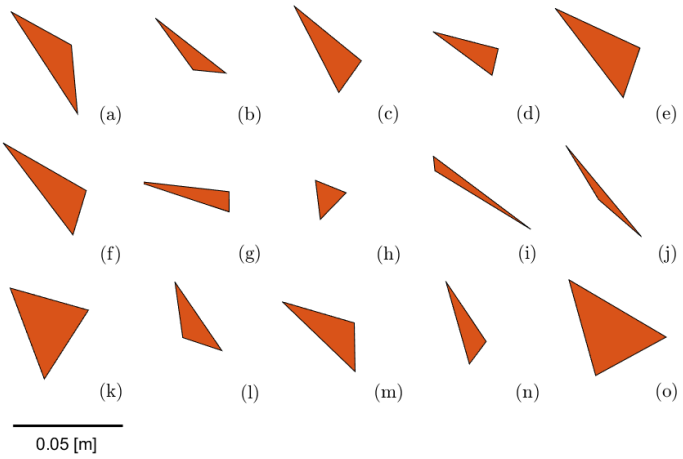

Figure 5.1: Randomly generated object shapes used in simulation experiments. Density of objects set to be the same as that of the simulated L-shape allen key.

We conducted three sets of experiments in simulation. Section 5.2.1 tests whether the Kruskal effect could be observed for the L-shaped object. Section 5.2.2 tested other triangular shapes to confirm that the effect is not specific to L-shaped objects. Finally, Section 5.2.3 introduced friction noise in our simulation, to observe its effect on convergence rate. Each set of experiments is described below.

5.2.1 Varying Random Action Sequences

The first set of simulation experiments used a single object, the L-shaped model of the allen key. We generated 43 distinct random action sequences S, each of length $N = 50$. Each sequence was repeated $M = 10,000$ times, starting from initial poses uniformly sampled from the CSpace as described above.

Figure 5.2 shows the entropy for each sequence, the mean across all sequences, and the interquartile range. In this instance, the Kruskal effect is readily observed. While the entropy is not monotonically decreasing, there is a clear trend. Of the 43 sequences tested, 29 converged

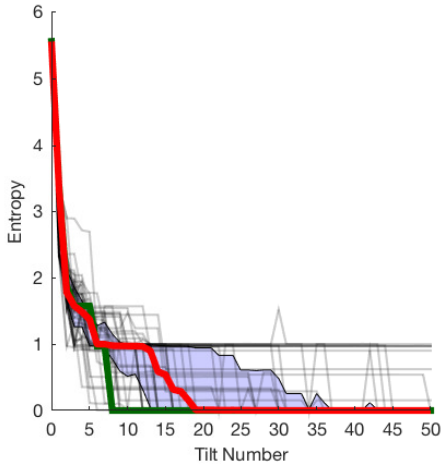

Figure 5.2: Kruskal effect for the allen key: $M = 10,000$ trials of $N = 50$ actions were repeated across 43 distinct random sequences. The mean (bold red line that converges by the 20^{th} tilt) is bounded by the interquartile range (in blue shaded region). The thin black lines show individual sequences' entropy trends, some of which approach zero by 50 actions. The shortest converging sequence is shown in green reaching zero entropy by the 8^{th} tilt.

to zero entropy, with all poses landing in a single voxel. The majority of data (25–75% or the interquartile range) is within the blue shaded region in Figure 5.2. On average, the entropy converges to a value close to zero by the 20^{th} tilt. The best randomly generated sequence converges in eight actions (shown as the green line in Figure 5.2), whereas the Erdmann and Mason plan, { `S, SW, SE, NE, SE` }, converges in five. While not conclusive, the results suggest that converging plans are common, but optimal plans are rare. This is expected since the actions were randomly chosen instead of being planned.

5.2.2 Varying the Object Shape

During the second set of experiments, we tested the effect of varying object shape. We used 15 different triangular object shapes (see Figure 5.1) and applied the fastest converging sequence we found for the allen key (shown as the green line in Figure 5.2). We conducted $M = 10,000$ repetitions for each object, starting at a randomly sampled initial object configuration.

The results are shown in Figure 5.3. Of the 15 objects tested, 10 converged to zero entropy.

The majority of the data shown by the interquartile range (blue shaded region in Figure 5.3) oriented the test object into a single final determined pose. On average, the entropy converges to a value close to zero by the 27^{th} tilt.

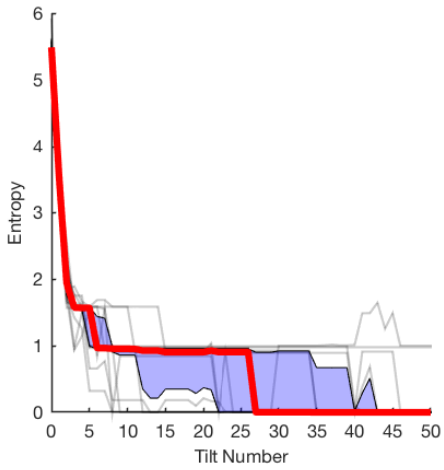

Figure 5.3: Varying the object shape: $M = 10,000$ trials of $N = 50$ actions were repeated using the same random sequence (green line from Figure 5.2) across 15 various object shapes. The mean (bold red line) is bounded by the interquartile range (in blue shaded region). The thin black lines represent distinct object poses, most of which converge by 50 actions.

The best sequence generated for the allen key does not perform as well on the other shapes, although it still tends to converge in most cases. One interpretation is that some objects are harder to orient than others, which is not surprising. In the context of pushing, this has already been proven [4]. Objects that are symmetrical or much larger in one dimension are harder to bring to zero entropy within 50 tilts (objects e, j, k, o, n in Figure 5.1). It is also likely that we have used a part-specific plan, by generating several part-agnostic plans and then selecting the best for the L-shaped object. We have only restricted to triangular shapes as an initial exploratory experiment, and studying other convex and concave object shapes is left for future work.

5.2.3 Varying the Friction Noise

The third set of experiments explored the effect of friction noise with the same 30-degree tray tilts and L-shaped object randomly initialized in the CSpace. We apply a simple noise model in which we let the coefficient of friction vary randomly with respect to position within the tray as our real experiments exhibited spatially varying friction due to wear. Figure 5.4 shows the low, medium, and high amplitudes of variation. We randomly generated 20 distinct friction maps, which were grouped by their mean friction to generate 13 low-noise maps, 3 medium-noise maps, and 4 high-noise maps. We used the allen key, and the best-performing action sequence found for the allen key in the first set of experiments, which is shown as the green line in Figure 5.2. The chosen action sequence that converges by eight tilts allows for observable medium and high friction noise convergence behavior, as low friction noise should converge close to eight tilts. We performed $M = 1,000$ repetitions for each friction map, which is sufficient to observe the effect on the probability distribution of pose between maps.

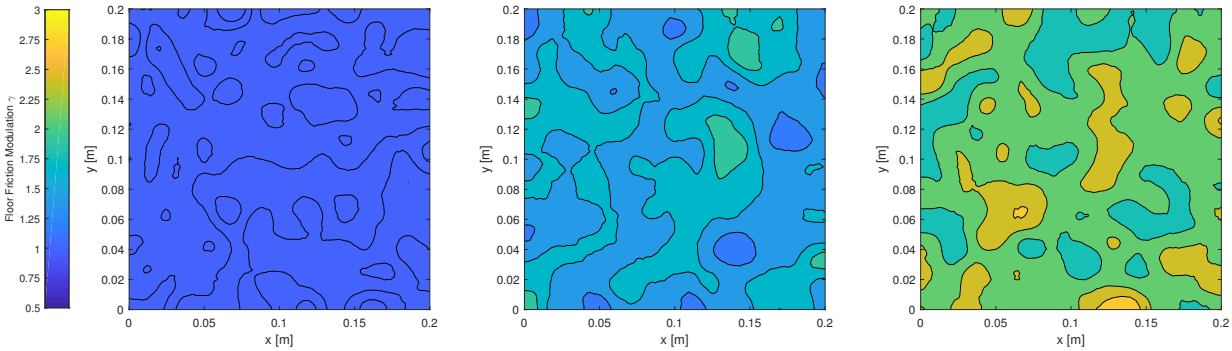

Figure 5.4: Spatially varying floor friction with low (left), medium (middle), and high (right) variation

The results are shown in Figure 5.5. For the low noise maps, 10 of the 13 entropy trends converged to zero entropy. On average, entropy converges to a value close to zero by the 9^{th} tilt.

Figure 5.5 also shows the results for medium and high noise. None of the medium- or high-noise maps converged to zero entropy. In both cases, the Kruskal effect is observable in that the general trend of entropy is decreasing, although they tend to not converge to zero entropy within 50 actions. These results show that lower friction noise positively affects the probability distri-

bution of object poses towards convergence. It is also likely that better-performing sequences exist for higher friction noise levels. Longer sequences are necessary to draw conclusions as to whether entropy for medium and high friction noise will level off or converge after more than 50 tilts.

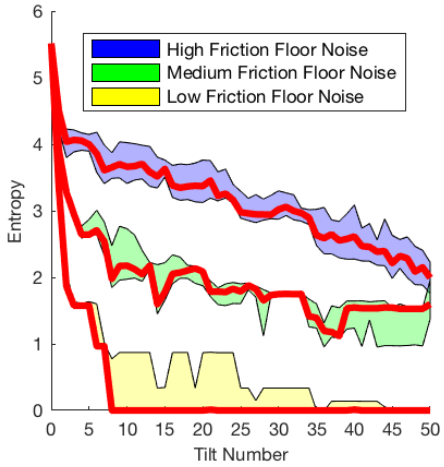

Figure 5.5: Varying the friction noise: $M = 1,000$ trials of $N = 50$ actions repeated using the same random sequence (green line from Figure 5.2) with 20 distinct floor friction noise amplitudes. The mean (bold red line) is bounded by the interquartile range for each category of friction floor noise illustrated in Figure 5.4—low, medium and high.

5.3 Physical Experiments

The simulation results suggests that the Kruskal effect can be observed for 2D objects, with significant entropy reduction for a variety of triangular objects and friction noise levels. However, physical rigid body interactions can be complex to simulate accurately. The goal of the physical experiments is to test the validity of the simulations. We tested one of the randomly generated sequences consisting of 50 actions. We used a 6-DOF ABB IRB 120 robotic arm tilting a 200 mm square aluminum tray. The object is a 77.5×27.5 mm allen key with an April Tag [17] to track the object with an overhead camera, as pictured in Figure 5.6. The tilting actions were

30-degree tilts in each of the eight cardinal directions. We ran M=500 trials which is less than the number of trials suggested by Equation 3.7, but due to wearing down of the tray from metal-metal interactions we restrict ourselves to fewer trials and thus lower resolution ($3 \times 3 \times 3$ grid). This inevitably leads to a less accurate estimate of the entropy, but we still expect to see the downward trend if the Kruskal effect is observed.

Figure 5.6: Experimental setup. An industrial robot tilts an allen key, with April Tag attached, in an aluminum tray. The overhead camera records the pose of the allen key after each tilt.

It is important that each trial be independent of the preceding trial, and that the initial poses approximate a uniform sampling of the CSpace. To that end, the robot shook the tray vigorously prior to the start of every sequence. The success of that approach is easily assessed by checking the initial entropy H^0. A uniform distribution over 27 voxels would yield an entropy of about 4.75. However, finite sampling from a uniform distribution is not likely to yield a uniform distribution. Numerical experiments for a dataset of 500 samples drawn into 27 bins suggested an expected initial entropy of approximately 4.72. The measured entropy of our initial distribution is around 3.9, for a difference of just under one bit. We attribute the difference to the fact that some of our CSpace volume $\mathcal{X} \times \mathcal{Y} \times \Theta$ is infeasible due to collisions with walls, and to small limitations in our vigorous shaking motion.

Entropy is calculated in the same way as Section 5.2, discretizing the tray volume into $3 \times$

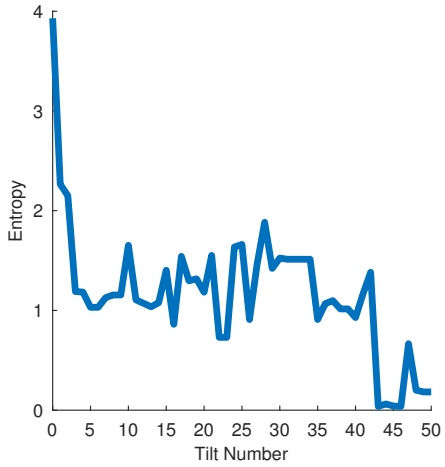

Figure 5.7: Robot Entropy Data: $M = 500$ trials of $N = 50$ actions repeated on the robot using the same sequence that was used in simulation experiments 5.2.2 and 5.2.3.

$3 \times 3 = 27$ (x, y, θ) voxels. Corresponding results are shown in Figure 5.7, where the lowest entropy of 0.0376 is achieved after the 43^{rd} tilt which is comparable to the final entropy achieved from executing the 5 optimal actions planned by Erdmann and Mason repeated over 500 trials as shown in Figure 5.8 [10]. The entropy line is quite noisy which makes it difficult to draw confident conclusions, but the general trend is downwards and indicative of the Kruskal effect. In future work, real world issues, such as wear and tear, should be addressed to obtain more trials and finer resolution for more concrete inferences.

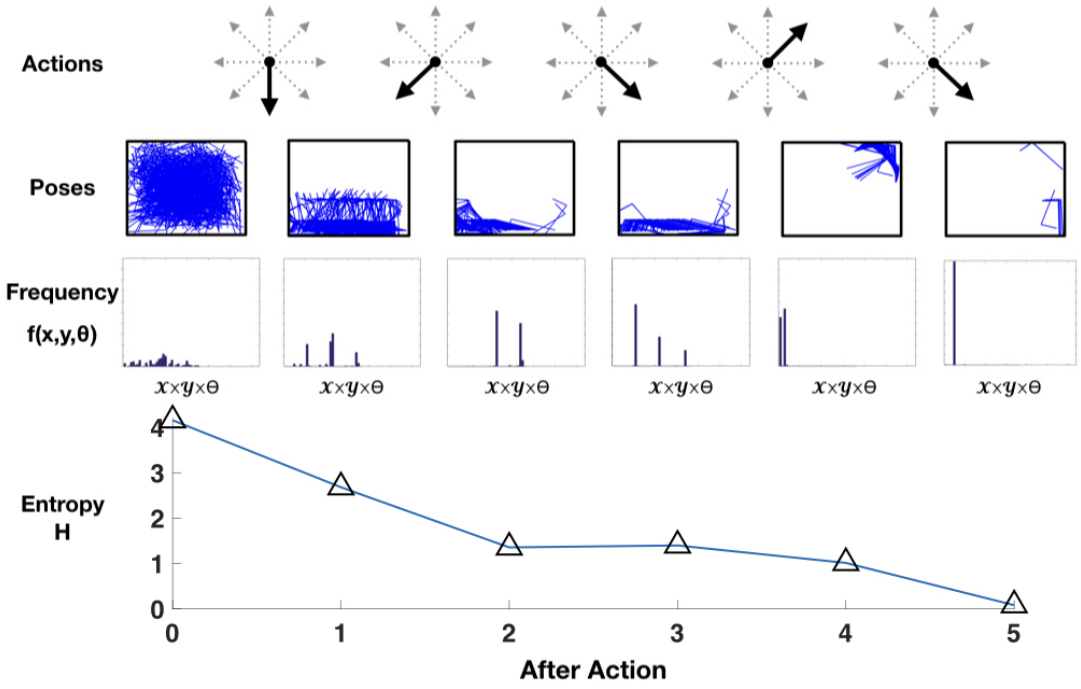

Figure 5.8: Robot Entropy with Planned Action Sequence: $M = 500$ trials of the optimal 5 actions to reorient an allen key repeated on the robot setup.

Chapter 6

Discussion

In this section, we will discuss the results presented in Sections 5.2 and 5.3, draw conclusions and discuss insights for future exploration.

From the planner proposed by Erdmann and Mason [10], we know that planned actions can orient an allen key to a final determined pose. Although their proposed sequence efficiently oriented the object, we wanted to explore how random sequences would perform at the same task. Towards this end, Section 5.2.1 tests various random sequences on the same test object. Almost all sufficiently long sequences significantly reduce the entropy, and most sequences result in zero entropy. We show that the Kruskal effect applies for any random sequence to mostly or completely reduce object pose uncertainty.

Given an object, it would be possible to produce an object-specific plan by searching random sequences and selecting the best. However, we consider action sequences that are not object-specific which is beneficial when introducing new objects. We show this in Section 5.2.2, where we selected the best allen key sequence, and repeated it for other triangular shapes. Figure 5.3 shows that the sequence reduced entropy to a few poses within 30 tilts. On average, the sequence succeeds at decreasing entropy for all tested objects, perhaps because the objects are all somewhat similar to the L-shaped object. Even a part-specific sequence serves as part-agnostic sequence, although a less efficient one. While we empirically only observed our results for square

trays and triangle shaped objects, the setup can be replicated with a freely-moving straight edge to push an object from any direction. The method of using randomized action sequences to reduce pose uncertainty is versatile. A possible extension of this work is to identify various objects and environments that are highly conducive to the Kruskal effect.

In Section 5.2.3, we explore the significance of non-deterministic actions, by introducing a noise model. While the entropy did generally decrease over the tilting sequence regardless of noise, the higher the noise, the slower the object poses seemed to converge. The Kruskal effect can be observed in less than ideal conditions such as high noise, but lower friction noises are more efficient at lowering pose uncertainty. Future work might extend the sequences to see if the entropy levels off at some value depending on the friction noise level or determine whether different sequences perform better at different noise levels.

In Section 5.3, we show that our theory can be applied to the real world. Even with the noise arising from variations in setup and execution, the object poses still converge to a relatively low entropy. In the future we are interested in further exploring the limits of tray tilting actions reducing object pose uncertainty in the real world and the effects of wear on physical systems through exploitation of material interactions.

Simulation provided large amounts of data and easily varied parameters to confirm the decrease in entropy provided by randomized action sequences. The largest entropy decrease among simulation and robot experiments was after the first move. At first, random initialization causes the object to be anywhere in the tray and after the first tile, the object rests only along the edges of the tray. Testing across different triangular object shapes demonstrated some of the generality in shapes that the system can tolerate. Simulation using different noise parameters showed that entropy reduction works under stochastic conditions.

In some of our sequences, the object pose did not converge completely to zero after 50 iterations. We think this is because for some objects, a certain mini-sequence of actions must be executed consecutively for distinct poses to converge to one pose. When randomly select-

ing actions, it may sometimes require a very long sequence for this mini-sequence to appear. This presents an opportunity to learn strategies that can predict these mini-sequences. Learned strategies can achieve high precision over a specific set of objects, and orthogonally our approach strives to achieve almost full convergence on a general set of objects. The majority of the entropy reduction occurs in the first 5 tilts. Figure 6.1 shows that the number of voxels occupied from there on continues to oscillate between 2 and 6 voxels. We can use learning to converge to a single pose more efficiently after the Kruskal effect achieves clusters or bi-modal states. This allows for less training data and higher accuracy in object pose. In essence, we can learn to find a mini-sequence of actions that efficiently reduces pose uncertainty.

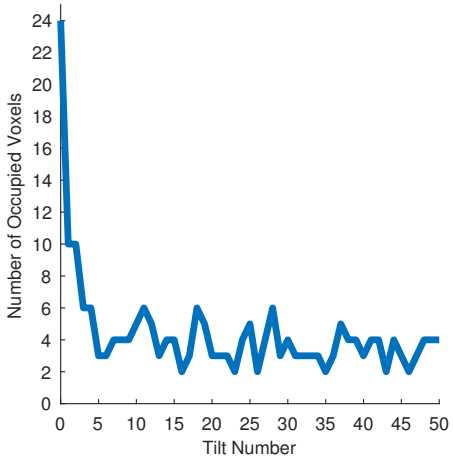

Figure 6.1: Robot Occupied Voxels Data: $M = 500$ trials of $N = 50$ actions repeated on the robot using sequence that was used in Section 5.3.

Additionally, for our experiments, small increases in entropy occur due to small changes between similar poses that map to distinct voxels. Later, these poses will converge again but may take some time to find the right actions to realign.

Informally, it is possible to make a few observations about the tray tilting process. The main order-producing phenomenon is when we drive the object pose to the boundary of the CSpace, i.e. a contact between object and tray wall. Ideally, this is a projection of the feasible poses to the boundary, and reduces the dimension of the feasible CSpace. For example, if each dimension

of $SE(2)$ is quantized into N bins, then at the beginning the pose is spread across N^3 bins, and after one action it has been projected to a surface spanned by N^2 bins.

In the simplest case, shown in Figure 3.1 using squeezing actions of a disk, this projection would be a normal projection onto a line. These actions would be analogous to tilting a tray back and forth. For a second action to combine most effectively with the first, the second line would be orthogonal to the first, and the final disk position would then be uniquely determined. As seen in Figure 3.1, starting from random initial poses of the disk, the pose uncertainty in (x,y) goes to zero if the squeeze grasps are orthogonal.

In general, the CSpace surfaces that correspond to kinematic constraints cannot be modeled as linear, nor are the projections linear, but still the toy example may provide some useful insights. The more closely two actions can be modeled as orthogonal projections, the better.

The main disorder-producing phenomenon might be sliding across the tray floor, where minor variations in friction can cause rotation of the object. The vagaries of sliding friction can also make it impossible to say whether an object will stick or slide along a tray wall.

There are also disorder-amplifying phenomena. For example, if the part strikes the wall sharply it will rebound, and the small variations in initial pose will be amplified over time to produce large variations. It is this effect we relied upon to randomize the object pose prior to testing a sequence of actions in our physical experiments.

The effectiveness of a sequence depends on how common and how effective the order-producing actions are, how frequently combinations occur, and how effectively they combine, versus the frequency and degree of disorder produced by the other actions. One goal of future work will be to explore this underlying structure more precisely, as a way of characterizing tasks.

Chapter 7

Conclusion

Examining the traditional approach of sense-plan-act, we observe the effects of an alternative approach of executing random sequences of actions without sensing. We show that a sufficiently long random sequence of actions can move an object from an unknown initial pose to a determined final pose, regardless of initial pose of the object, varying object shapes, and stochasticity in the environment. This effect is explored in greater detail through simulation using millions of tilts and observing the entropy trends over action sequences. We learned how some parameters affect our system: longer sequences lower object uncertainty more, and stochasticity in the environment as well as some variation in triangular object shapes does not disturb the system. We also illustrated the same effect on a real robot and saw a decreasing trend in entropy. However, the final entropy is not as low as suggested by simulation results, due to real world challenges and complications such as wear and tear.

This is a different paradigm than the sense-plan-act approach where the final pose and the action sequence to achieve that pose are planned; exploring this alternative paradigm and its limitations could be fruitful. We offer insights into the idea of randomized action sequences instead of planning. The advantage in our setup is that random tray-tilting actions are not part-specific and reduce system complexity for new objects. The Kruskal effect can be useful for orienting kit of parts, a difficult planning problem often encountered in industrial applications.

Kit of parts reorienting involves placing similarly shaped objects in a compartmentalized tray where each compartment holds one of these objects. A naive planning approach for reorienting would require us to do joint planning for all the objects in the tray whose complexity would scale with the number of objects. Using the proposed approach, we can orient many objects without scaling complexity by the number of objects. More specifically, we can execute the proposed approach on such a compartmentalized tray and by the end of a random sequence, we expect most objects to be in a determined pose. The manipulator can retrieve each object from the fixed pose and continue to the next task. If the object was not successfully grasped, it can benefit from more tilts as the rest of the compartments are refilled and the sequence of random actions continues. The more tilts an object undergoes, the more certain its final position will be. Thus, many similarly shaped objects can be simultaneously oriented for a subsequent task.

Finally, the proposed approach is not limited to tray tilting. A randomized sequence of actions can come in the form of any set of actions performed by a manipulator. For our purposes, this was a tray attached to the end of an industrial arm. An example of an alternative application of this work can be to replace the tray with a freely mobile wall, where the set of possible actions is pushing an object from any angle. In such a robotic system, we would expect the object to reach a determined pose after a random sequence of actions as well, exhibiting the Kruskal effect.

Chapter 8

Future Work

Extensions of this work to various polygons or approximations of non-convex objects and tray-tilting alternatives, such as pushing, would further explore the effects of randomized action sequences. The sustainability of our approach can be tested through longer sequences in simulation and on physical systems, as well as more trials for higher quality estimates of entropy. We are also interested in ways to capture the order of the system in a data-efficient way. A future goal is to move towards a tray with a lid that can offer a 3D exploration of part-agnostic tray-tilting to determine 3D object pose.

To identify action sequences that are efficient at orienting a given object, we could learn a policy, such as that by Christiansen et al. [8], but with finer discretized tray regions for more accurate object poses. Another future direction would be to explore potential applications of the proposed approach in simplifying a pose estimation problem for a manipulation task.

Bibliography

[1] American Mathematical Society The Kruskal Count. 2017. URL http://www.ams.org/samplings/feature-column/fcarc-mulcahy6. 1

[2] Srinivas Akella and Matthew T. Mason. Orienting Toleranced Polygonal Parts. *The International Journal of Robotics Research*, 19(12):1147–1170, 2000. doi: 10.1177/02783640022068002. URL https://doi.org/10.1177/02783640022068002. 2

[3] R. P. Berretty, K. Goldberg, M. H. Overmars, and A. F. van der Stappen. Orienting parts by inside-out pulling. 2:1053–1058 vol.2, 2001. ISSN 1050-4729. doi: 10.1109/ROBOT.2001.932733. 2

[4] Robert-Paul Berretty, Ken Goldberg, Mark H. Overmars, and A. Frank van der Stappen. On fence design and the complexity of push plans for orienting parts. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, SCG '97, pages 21–29, New York, NY, USA, 1997. ACM. ISBN 0-89791-878-9. doi: 10.1145/262839.262853. URL http://doi.acm.org/10.1145/262839.262853. 5.2.2

[5] Karl Böhringer, Russell Brown, Bruce Donald, Jim Jennings, and Daniela Rus. Distributed robotic manipulation: Experiments in minimalism. In Oussama Khatib and J. Salisbury, editors, *Experimental Robotics IV*, volume 223 of *Lecture Notes in Control and Information Sciences*, pages 11–25. Springer Berlin / Heidelberg, 1997. ISBN 978-3-540-76133-4. 10.1007/BFb0035193. 1, 2

[6] Randy C. Brost. Automatic Grasp Planning in the Presence of Uncertainty. *The International Journal of Robotics Research*, 7(1):3–17, 1988. doi: 10.1177/027836498800700101. 2

[7] Gregory S Chirikjian. *Stochastic Models, Information Theory, and Lie Groups, Volume 1: Classical Results and Geometric Methods*, volume 1. Springer Science & Business Media, 2009. 2, 3

[8] A. D. Christiansen, M. T. Mason, and T. M. Mitchell. Learning reliable manipulation strategies without initial physical models. *Proceedings., IEEE International Conference on Robotics and Automation*, pages 1224–1230 vol.2, May 1990. doi: 10.1109/ROBOT.1990.126165. 8

[9] Nikhil Chavan Dafle and Alberto Rodriguez. Prehensile pushing: In-hand manipulation with push-primitives. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, September 2015. 2

[10] M. A. Erdmann and M. T. Mason. An exploration of sensorless manipulation. *IEEE Journal on Robotics and Automation*, 4(4):369–379, Aug 1988. ISSN 0882-4967. doi: 10.1109/56.800. 1, 2, 5.1, 5.2, 5.3, 6

[11] Kenneth Y. Goldberg. Orienting polygonal parts without sensors. *Algorithmica*, 10(2): 201–225, Oct 1993. ISSN 1432-0541. doi: 10.1007/BF01891840. URL `https://doi.org/10.1007/BF01891840`. 2

[12] D. D. Grossman and M. W. Blasgen. Orienting mechanical parts by computer-controlled manipulator. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-5(5):561–565, Sept 1975. ISSN 0018-9472. doi: 10.1109/TSMC.1975.5408381. 2

[13] David M. Lane. *Online Statistics Education*. Rice University. URL `http://onlinestatbook.com/`. 3

[14] Kiju Lee, Matt Moses, and Gregory S. Chirikjian. Robotic Self-replication in Structured

Environments: Physical Demonstrations and Complexity Measures. *The International Journal of Robotics Research*, 27(3-4):387–401, 2008. doi: 10.1177/0278364907084982. 3

[15] Steve Miller. Simscape Multibody Contact Forces Library, 2017. URL `https://www.mathworks.com/matlabcentral/fileexchange/47417`. 5.2

[16] Robert Paolini Bowei Tang Siddhartha S. Srinivasa Michael Erdmann Matthew T. Mason Ivan Lundberg Harald Staab Nikhil Chavan Dafle, Alberto Rodriguez and Thomas Fuhlbrigge. Extrinsic dexterity: In-hand manipulation with external forces. *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1578–1585, May 2014. 2

[17] Edwin Olson. AprilTag: A robust and flexible visual fiducial system. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3400–3407. IEEE, May 2011. 5.3

[18] A. Sanderson. Parts entropy methods for robotic assembly system design. *Proceedings. 1984 IEEE International Conference on Robotics and Automation*, 1:600–608, Mar 1984. doi: 10.1109/ROBOT.1984.1087155. 2, 3

[19] P. L. Várkonyi. Estimating part pose statistics with application to industrial parts feeding and shape design: New metrics, algorithms, simulation experiments and datasets. *IEEE Transactions on Automation Science and Engineering*, 11(3):658–667, July 2014. ISSN 1545-5955. doi: 10.1109/TASE.2014.2318831. 2

[20] J. Zhou, R. Paolini, A. M. Johnson, J. A. Bagnell, and M. T. Mason. A probabilistic planning framework for planar grasping under uncertainty. *IEEE Robotics and Automation Letters*, 2(4):2111–2118, Oct 2017. doi: 10.1109/LRA.2017.2720845. 2